

The Emperor's New 'Knows'

KENT BACH

When I examine contextualism there is much that I can doubt. I can doubt whether it is a cogent theory that I am examining, and not a cleverly stated piece of whacks. I can doubt whether there is any real theory there at all. Perhaps what I took to be a theory was really some reflections; perhaps I am even the victim of some cognitive hallucination. One thing however I cannot doubt: that there exists a widely read patch of a round and somewhat bulgy shape.

- a traditional epistemologist¹

The title of this paper calls for it to stick to the obvious. Even if it did, it would probably not convince the contextualist. Knowing that, I will be comforted by the thought that whether or not 'knows' is a context-sensitive term, at least 'obvious' and 'convincing' are. Perhaps 'context-sensitive' is context-sensitive too.

I begin, in Section I, with what contextualism says, what it doesn't say, and what it implies about knowledge attributions. Even if contextualism is true and, contrary to invariantism, a given knowledge-ascribing sentence can express various propositions in various contexts, those propositions are not themselves context-bound. This is something that contextualists do not make clear. In section II, I will sketch the contextualist's strategy for containing skepticism and discuss whether this strategy really explains why unsuspecting people can be duped by skeptical arguments. An alternative explanation is that the conflicting intuitions that give rise to skeptical paradoxes don't really bear on the

¹ The allusion here, if it needs to be made explicit, is to this famous passage from traditional epistemologist H. H. Price's *Perception* (1932, 3): "When I see a tomato there is much that I can doubt. I can doubt whether it is a tomato that I am seeing, and not a cleverly painted piece of wax. I can doubt whether there is any material thing there at all. Perhaps what I took to be a tomato was really a reflection; perhaps I am even the victim of some hallucination. One thing however I cannot doubt: that there exists a red patch of a round and somewhat bulgy shape, standing out from a background of other colour-patches, and having a certain visual depth, and that this whole field of colour is of colour is directly present to my consciousness."

truth conditions of knowledge attributions but are merely vacillating responses to skeptical considerations. In any case, as claimed in section III, contextualism doesn't really come to grips with skepticism. In attempting to confine the plausibility of skeptical arguments to contexts in which far-fetched skeptical possibilities are raised, it concedes both too much to the skeptic and too little. Also, as section IV points out, in arguing against invariantism contextualists have mainly focused on the skeptical variety, according to which knowledge requires the highest degree of evidence, justification, and conviction. Although the contextualist objections to skeptical invariantism are not cogent, this view is independently implausible.

Much more plausible is moderate (nonskeptical) invariantism, a version of which I will propose in section V. From its perspective, the evidence that seems to support contextualism appears in a very different light. In contexts where special concerns arise, whether skeptical or practical, what varies is not the truth conditions of knowledge attributions but the knowledge attributions people are prepared to make. It is not the standards for the truth of knowledge attributions that go up but the attributor's threshold of confidence regarding the relevant proposition. When that happens, as in the examples contextualists rely on, people require stronger evidence than is necessary for knowing. That's what it takes for them to eliminate residual doubts and to attribute knowledge to others. So my version of moderate invariantism is a kind of error theory, but not an extreme error theory like contextualism and skeptical invariantism.

Finally, as I will suggest in section VI, part of what makes a belief justified is that the cognitive processes whereby it is formed and sustained are sensitive to realistic counterpossibilities (so-called relevant alternatives). The very occurrence of the thought of a counterpossibility gives one *prima facie* reason to take it seriously, and the fact that a counterpossibility does not come to mind is evidence for its irrelevance. But that fact is evidence that one cannot explicitly consider, since to do so would be to bring the counterpossibility to mind. Examining this underappreciated phenomenon will shed new light on why possibilities that are irrelevant to knowing are properly ignored.

Before we get down to business, a parable is in order.

The Dirtmatist and the Septic

The Dirtmatist thinks that he can keep his hands clean by washing them with a little soap and water every so often. The Septic thinks that because germs are everywhere, it is impossible to keep his hands clean. The best he can do is to scrub his hands repeatedly with industrial-strength cleaning agents and hope for the best. One day the Dirtmatist encounters the Septic near the sink in the men's room, and offers to shake hands. The Septic backs off in fear. The Dirtmatist assures him, "My hands are clean." The Septic retorts, "No they're not," and backs this up with the following argument:

Septical Argument

If your hands were clean, they would be free of contaminants.

Your hands are not free of contaminants.

Your hands are not clean.

At least the Septic doesn't doubt the existence of the Dirtmatist's hands — it's only their dirtiness that worries him. Anyhow, the Dirtmatist doesn't buy the Septic's argument. "I've just washed my hands," he protests. Even though he's not at all naive about the microscopic world, he rejects the second premise. He just doesn't worry about germs or dirt particles too small for the eye to see. For him, it's out of sight, out of mind. For the Septic, it's out of mind, still in body.

The Dirtmatist and the Septic argue for a while, until in walks Notsick, a more sophisticated thinker. He accepts the truth of the second premise but rejects the first, which is supported by what he refers to as the Cleanser Principle. At this point the Dirtmatist and the Septic join sides (not that they shake hands), both thinking that Notsick is being too clever by half. Despite their disagreement about the second premise, they find it unpalatable to reject the Cleanser Principle. They gang up on Notsick, but he sticks to his guns. Then they get back to arguing with each other. Finally, a Cleantextualist emerges from a stall and comes to the rescue.

After washing his hands, the Cleantextualist assures the Dirtmatist that he was right when he uttered, "My hands are clean." He concedes to the Septic that, yes, there is no way to eliminate every last germ and particle of dirt. And, while acknowledging Notsick's noble antiseptic intentions, he chides him for rejecting the axiomatic Cleanser Principle. And though it might seem that the Cleantextualist all but concedes the Septical

Argument, he hasn't really. He points out something overlooked all along by the others, that 'clean' is context-sensitive and that 'contaminant' is too. It turns out, much to everyone else's surprise, that what the Dirtmatist asserted is not what the Septic argued against.

I What Contextualism Says and Implies

Fred Dretske expresses the natural intuition that "factual knowledge is absolute. It is like being pregnant: an all or nothing affair" (1981: 363). One can be newly pregnant but not a little pregnant, or almost ready to deliver but not highly pregnant. Similarly, as Dretske observes, "I can have a better justification than you, but my justification cannot be more sufficient than yours." It can be more than sufficient for knowledge, but not more sufficient for knowledge. Justification (and evidence) comes in degrees, but knowledge does not. Now does contextualism conflict with any of this? Not at all. Contextualists say that what is sufficient for knowledge varies with the context in which knowledge is attributed. As we will see, however, that is not quite what they mean.

One can take a contextualist position about various expressions, such as 'obvious', 'tall', and 'good', as well as 'knows'. Contextualism about a given expression (or class of expressions) is a semantic thesis. It says that any sentence containing the expression, even if otherwise free of ambiguity, indexicality, and vagueness (or if the effects of these are kept fixed), expresses different propositions (or, if you prefer, has different truth conditions) in different contexts of utterance.² Here, since we are here concerned solely with contextualism about 'knows' and knowledge-ascribing sentences, I will use the label 'contextualism' specifically for epistemic contextualism.

² This sort of contextualism, which concerns specific expressions, is not to be confused with the kind that prevails in some philosophy of language circles. There the term 'contextualism' is used for a rather radical family of theses about sentence meaning, such as that not just a great many but virtually all sentences do not express complete propositions, that pragmatics intrudes into semantics in the sense that "what is said" is generally determined partly by pragmatic factors, and that the meanings of a great many lexical items are semantically impoverished and require contextual enrichment. In 'Context ex Machina' (Bach forthcoming), I suggest that the platitudes that motivate such theses are misstated or overstated. When these phenomena are accurately characterized, by taking certain independently motivated distinctions into account, the motivation for such theses loses its force. The simplest distinction to observe is that between content being determined *by* context and content being determined *in* context (but by something else). Disregarding this distinction tends to lead contextualists, as well as many of their critics, to use phrases like 'context-dependent' and 'context-sensitive' interchangeably with 'contextually variable', and then to treat the relevant phenomena as having semantic import. Epistemic contextualists tend to do likewise.

Contextualism directly concerns knowledge attributions, not knowledge. In fact, it is a thesis about the semantic contents of knowledge-ascribing sentences, not just what people implicate or presuppose when uttering them.³ It claims that a sentence of the form ‘S knows (at t) that p’ can be true as uttered in one context and false as uttered in another, depending on the epistemic standards that govern the context. The claim is not merely that people’s willingness to make a given knowledge attribution depends on the standards but that the standards governing the context actually affect which proposition the knowledge-ascribing sentence expresses in that context.

It is crucial to see, although contextualists do not stress this, that contextualism does not imply that the proposition expressed by a given knowledge-ascribing sentence in a given context can itself have different truth values in different contexts. Contextualism does not imply that somebody can know something if the attributor’s standards are low and fail to know it if they are high. Nor does it imply that somebody can both know something relative to one context of attribution, and not know it relative to another. What it does imply is that a sentence of the form ‘S knows (at t) that p’ can be true as uttered in one context and false as uttered in another. This is not because the proposition the sentence expresses has a different truth value, but because the sentence expresses a different proposition. That is something contextualists recognize but, it seems, do not always keep in mind.⁴

³ On what people implicate and presuppose, see Rysiew forthcoming. It is important to keep in mind that what have semantic contents are sentences, not utterances. That is why David Kaplan distinguishes a “sentence-in-a-context” from an utterance of the sentence (1989: 522). This distinction is essential to my formulation of the semantic-pragmatic distinction (Bach 1999). The basic idea is that information counts as pragmatic if it derives not from the content of the sentence but from the fact that the sentence is actually uttered.

⁴ As Stewart Cohen explains, “strictly speaking, instead of saying that S knows in one context [of attribution, not S’s context] and fails to know in another, one should really say that ‘S knows that P’ is true in one context and false in the other” (1999: 65). Rather than use metalinguistic locutions, he prefers the less “stylistically cumbersome” object language but advises the reader not to be misled by this. Still, it is easy to mislead the reader, as when he says, for example, “the standards that determine how good one’s reasons have to be in order to know are determined by the context of ascription” (1999: 59). Lewis (1996) makes no bones about misleading the reader. It is not until his very last paragraph, after making a brilliant series of startling and sometimes paradoxical observations and suggestions about knowledge and the knowledge-destroying effect of epistemology, that he acknowledges, “I could have said my say fair and square, breaking no rules. It would have been tiresome, but it could have been done. The secret would have been to resort to ‘semantic ascent’” (1996: 566).

Contextualists do make clear that the context they have in mind is not the epistemic context of the subject of the knowledge attribution. Everybody agrees that what it takes for George to know that he has hands or, to put it more accurately, for the sentence ‘George knows that he has hands’ to be true, can depend on George’s epistemic situation. This is a matter not of setting standards but of meeting them. Obviously how hard it is to know something does not depend just on the thing to be known but also on the situation of the prospective knower. If there are considerations that need to be taken into account (one’s memory has been shaky lately), possibilities to consider that ordinarily can be ignored (maybe the zoo keepers have placed a cleverly painted mule in the zebra cage), or alternatives to eliminate (a person’s twin has returned), then the subject must reckon with them. This may be because of things the subject is aware of (or at least has reason to suspect) or facts about his circumstances that he needs to be aware of. If Austin (1961) was right, possibilities or alternatives are relevant only if there are special reasons to consider them. So the subject’s context, insofar as it affects his epistemic position, can bear on the truth of a knowledge attribution. But contextualism concerns the attributor’s context, which can vary even while the subject’s epistemic position stays fixed, and claims that this context bears on the content of the attribution.⁵

How can it be that a sentence like ‘George knows that he has hands’, even with time and references fixed, does not have a fixed propositional content? Doesn’t the verb ‘knows’ express an invariant two-term relation between the knower and the known?⁶ Contextualists tend to be not all that clear about this. They don’t claim that ‘know’ is

⁵ See Heller 1999a for an especially clear explanation of how (from a contextualist perspective) this can be. I should add that in the case of first-person knowledge attributions, where the subject and the attributor are one and the same, it might seem puzzling (even from a contextualist perspective) how the subject’s epistemic position can remain fixed while the content of a self-attribution of knowledge can vary. However, the standards for evaluating such an attribution, even if dependent, say, on the intentions of the (self-) attributor, can vary, for reasons independent of that person’s, qua subject, epistemic position. Still, contextualists should not to focus as much as they do on first-person cases. Focusing on cases in which attributor and subject are one and the same can only muddy the waters.

⁶ The terms ‘contextualism’ and ‘invariantism’ were coined by Peter Unger (1984: 6-11). Arguing that there is a trade-off between their respective virtues and vices, he concludes that there is no fact of the matter as to whether contextualism or invariantism is correct. He adopts this position of “semantic relativity” not just on ‘know’ but also on gradable terms that can seem to be absolute, such as ‘flat’ and ‘empty’.

ambiguous,⁷ but some suggest that it is context-sensitive because it is a kind of indexical (Cohen 1988) and others because it is vague (Heller 1999a).⁸ Some are reluctant to commit themselves as to its semantic character (indexical, vague, or something else) and are content to say that the “standards” for knowing, or what “counts as” knowing (DeRose 1995), depends on the context, or on what possibilities are “properly ignored” (Lewis 1996).⁹ There are some delicate issues here — these are not matters of incidental detail — but I will not be addressing them.¹⁰

Regardless of its detailed formulation, contextualism entails either that ‘know’ expresses different relations in different contexts or that it expresses a single relation that is relativized to a contextually variable epistemic standard.¹¹ Either way, ‘know’ has variable content. It is incoherent to suppose that it expresses a single, unrelativized relation and yet that identical knowledge attributions made in different contexts can differ in truth value. Contextualists cannot coherently mean, even if they often say, that the

⁷ Of course it has an acquaintance sense, corresponding to the French ‘connaître’ and the German ‘kennen’ as opposed to ‘savoir’ and ‘wissen’, but we are ignoring that sense and limiting our attention to ‘know’ as followed by a clause.

⁸ Noting that “the penumbras of vague terms can dilate or constrict according to conversational purposes,” Schiffer points out that if the context variability of ‘know’ consisted simply in its vagueness, this sort of variability would be “of no use to the contextualist, [because] speakers are perfectly aware of it when it’s going on” (1996: 327-8).

⁹ Unger (1986: 130-1) lists assorted factors, involving the subject’s psychological state, his justification, and what he can rule out, as well as rationality, reliability, and possibility.

¹⁰ Two issues are worth noting. Contextualists sometimes seem to suppose that what changes the standards is the salience of improbable or even far-fetched possibilities. However, if such a change is supposed to affect the semantic content of a ‘knows’-ascription, salience cannot be what affects it. Salience is obviously a feature relevant to pragmatics, not to semantics (see Bach 1999 and forthcoming). It plays a role in what speakers are likely to mean when they say what they say. Rysiew 2001 develops a plausible account of its pragmatic role in knowledge attributions. Also, there are linguistic issues to contend with. As Jason Stanley (forthcoming) argues, ‘know’ does not behave like ordinary indexicals (‘I’, ‘tomorrow’), relational terms (‘local’, ‘enemy’), or gradable adjectives (‘tall’, ‘flat’). For a probing semantic analysis of such adjectives and comparison of relative (‘tall’, ‘rich’) with absolute adjectives (‘flat’, ‘empty’), see Kennedy forthcoming, and for an ingenious semantic- pragmatic account of how absolute terms work see Lasersohn 1999. It is curious that Cohen, who argues that the context-sensitivity of ‘knows’ derives from that of ‘justified’, likens the relative term ‘justified’ to the absolute term ‘flat’ rather than to a relative term like ‘tall’.

¹¹ Contextualists differ as to whether epistemic standards are a matter of degree of justification, extent of relevant alternatives, or range of possible worlds in which the truth is tracked. I’ll ignore this difference here. Also, insofar as they distinguish standards simply by their strength, they implicitly and implausibly assume that standards form a linear ordering.

standards for knowledge or what counts as knowing can vary with the context.¹² It is somewhat better to say what it takes for a given knowledge-ascribing sentence to be true can so vary, but this must be understood to mean that its truth value can vary only because its content can vary. The same content cannot be true in one context and false in another. Stewart Cohen is clear on this:¹³

How from the viewpoint of formal semantics should we think of this context-sensitivity of knowledge ascriptions? We could think of it as a kind of indexicality. On this way of construing the semantics, ascriptions of knowledge involve an indexical reference to standards. So the knowledge predicate will express different relations (corresponding to different standards) in different contexts. But we could instead view the knowledge predicate as expressing the same relation in every context. On this model, we view the context as determining a standard at which the proposition involving the knowledge relation gets evaluated. So we could think of knowledge as a three-place relation between a person, a proposition, and a standard. (1999: 61)

As Cohen recognizes, “As long as we allow for contextually determined standards, it doesn't matter how formally we construe the context-sensitivity. These semantic issues, as near as I can tell, are irrelevant to the epistemological issues.” Using ‘D’ to represent the standard determined by context, we can capture the contextualist conception of the variable content of a simple knowledge-ascribing sentence by means of a more elaborate one that makes the relevant standard explicit. We can do this in either of two ways:

indexed: ‘S knows_D at t that p’

relativized: ‘S knows at t relative to D that p’.

¹² Here are some examples of what they say: “One speaker may attribute knowledge to a subject while another speaker denies knowledge to the same subject, without contradiction” (Cohen 1988: 97); “In some conversational situations, one’s epistemic position must be stronger than in others to count as knowing” (De Rose 1995: 30); “What counts as having this property [e.g., of knowing that grass is green] might vary from context to context” (Kompa 2002: 88). Such ways of putting things misleadingly suggest that the truth value of a knowledge attribution can somehow vary with context while its content remains fixed.

¹³ He is not so clear on his argument for contextualism: “Justification, or having good reasons, is a component of knowledge, and justification certainly comes in degrees. So context will determine how justified a belief must be in order to be justified *simpliciter*. This suggests a further argument for the truth of the contextualist’s claim about knowledge. Since justification is a component of knowledge, an ascription of knowledge involves an ascription of justification” (1999: 60). This is a weak argument. As Richard Feldman points out, “from the fact that the word ‘justified’ displays context sensitivity, it does not follow that the necessary condition for knowledge is similarly context sensitive. ... It could be that the degree of justification needed for knowledge is unchanging” (2001: 67). Not only is it entirely compatible with Cohen’s assumptions that knowledge requires a certain fixed degree of justification, this degree could be the highest degree of justification. Stanley (forthcoming, sec. 3) offers more complicated objections to Cohen’s argument.

The effect is the same either way:¹⁴ a sentence of the form ‘S knows at t that p’ does not express a complete proposition except relative to a standard, and the standard is determined (somehow) by the context.¹⁵ Either way, ‘knows’ does not express a fixed two-term relation. It expresses either a contextually variable two-term relation or a fixed three-term relation whose third term, the operative standard, varies with context. And, as contextualists stress, “there is no context independent correct standard” (Cohen 1999: 59). But it must also be stressed that no matter how context “determines” the standard that figures in the content of a knowledge-ascribing sentence, the *content* is not hostage to the context. This content is a proposition that can be expressed in a context-independent way by means of a more elaborate knowledge-ascribing sentence that makes the relevant standard explicit, either indexed (‘S knows_D at t that p’) or relativized (‘S knows at t relative to D that p’). So even if which proposition a simple knowledge-ascribing sentence depends on the context, the proposition thus expressed is context-independent.

Accordingly, in order to indicate that the word ‘know’ does not express a fixed two-term relation, from now on, at least in a contextualist context, I will put it in brackets and say that someone [knows] something. In such a context it would be better to call knowledge-ascribing sentences ‘*knows*’-*ascriptions* and to call assertive utterances of such a sentence [*knowledge*] *attributions*. [*Knowledge*] *denials* are assertive utterances of the negation of such a sentence (‘S does not know that p’).

Consider the effect for contextualism if some such device is not used. How would someone in one context report (or believe) a knowledge attribution made by someone

¹⁴ As Jonathan Schaffer has reminded me, their effects are not the same in special linguistic environments, such as in ellipsis and in focus constructions. As he argues in section 3 of Schaffer forthcoming, the relativized approach is truer to the data; he concludes that ‘knows’ expresses a ternary relation and is not an indexical.

¹⁵ Cohen goes on to ask, “How precisely do the standards for these predicates get determined in a particular context of ascription? This is a very difficult question to answer. But we can say this much. The standards are determined by some complicated function of speaker intentions, listener expectations, presuppositions of the conversation, salience relations, etc. — by what David Lewis calls the conversational score” (1999, 61). He does not explain how such seemingly pragmatic factors can contribute to semantic content. Neither does DeRose (forthcoming), who takes the determination of standards to be a matter of implicit negotiation. Here he relies on a distinction between the “personally indicated” standards of the individual participants and the standards that actually contribute to the truth conditions of a knowledge attribution at a given stage in a conversation.

else in another context where the prevailing standards are different? For example, if Martha said, “George knows that he has hands,” and you later report this with (1),

(1) Martha said that George knows that he has hands.

then according to contextualism your use of ‘knows’ should be sensitive to your context, not Martha’s. But this means that in uttering (1), you are not reporting what Martha said. Indeed, as Nikola Kompa (2002: 83) points out, contextualism predicts that you could say something true in uttering (2):¹⁶

(2) Martha said something true in uttering ‘George knows that he has hands’, but George does not know that he has hands.

This “unpleasant consequence” of contextualism, as Kompa calls it, can be avoided only if the relevant standards are made explicit, as in (3), or at least if there is some indication that the standards are different, as in (4):

(3) Martha said something true in uttering ‘George knows [relative to D_2] that he has hands’, but George does not know relative to D_1 that he has hands.

(4) Martha said something true in uttering ‘George knows [relative to some standard distinct from D_1] that he has hands’, but George does not know relative to D_1 that he has hands.

So the contextualist is faced with the problem of explaining how it is that we can use sentences like (1), which makes no mention of standards, to report what someone says (or thinks) someone else knows.

A contextualist would not respond by insisting that shifts in standards occur only when epistemologists raise skeptical possibilities and that otherwise epistemic standards stay fixed. As Keith DeRose explains (1999: 195), an essential part of the case for contextualism is that standards are sometimes raised in everyday contexts, not radically but still substantially. Supposedly this is what happens in DeRose’s (1992: 913) and Cohen’s (1999: 58) well-known Bank and Airport examples. Contextualists rely on such examples to show that “our ordinary intuitions” are responsive to alleged variations in the contents of [knowledge] attributions. So they do need to confront the problem posed by reporting on what someone says or thinks someone else knows, especially when, as illustrated by (2) above, the reporter’s context is the stronger. In the case of a report of a

¹⁶ Cappelen and Lepore (2003) thoroughly develop this very point.

[knowledge] denial, the problem is clearest when the reporter's context is the weaker, as in this variant of (2), "Martha said something true in uttering 'George does not know that he has hands', but George does know that he has hands."

What does contextualism predict if you encounter a [knowledge] attribution out of context? It seems to predict that you won't be in a position to grasp which proposition the sentence expresses. Suppose you eavesdrop on the middle of a conversation and hear one person say to the other, "Nixon knew that Liddy was planning the Watergate break-in." Since it is not evident to you which [knowledge] relation 'knew' expresses, you can have only a vague idea of what is being said. Lacking any specific information about the context in which the [knowledge] attribution was made, you should feel a bit uncertain as to what was said. But you won't. So far as I can tell, to avoid this difficulty the contextualist would have to show that there is some unique default [knowledge] relation that people presumptively take to be expressed by 'knows'. This approach would be implausible for 'flat' or 'tall', but maybe it could work for 'knows'.

For what it's worth, notice that explicitly relativized knowledge attributions and denials sound rather strange:

(5) ?Jack knows relative to ordinary standards that there's water at the top of the hill.

(6) ?Jill doesn't know relative to high standards that there's water at the top of the hill.

Comparative and degree-modified knowledge attributions sound strange too:

(7) ?Jack knows relative to a higher standard than Jill does that the hill is steep.

(8) ?Jill knows very highly/strongly that Jack fell down.

(9) ?Jack somewhat/nearly/barely knows that Jill tumbled down the hill.

(10) ?Jill knows to a high/some degree that she should have stayed home.

It is not clear what to make of the marginal status of such sentences.¹⁷ Perhaps these sentences sound bad only because language users are not imbued with the insights of contextualism. If people were cognizant of the context variability of 'knows' and the various relations it expresses, or at least realized that knowledge is standards-relative, then maybe such forms would not only sound all right but would be in common use. As things are, however, "no ordinary person who utters 'I know that p', however articulate,

¹⁷ For discussion of whether and in what ways 'knows' is gradable, see Stanley forthcoming, sec. 2, and Ludlow this volume. Stanley points out that a sentence like this variant of (8) is all right, 'Jill knows very well that Jack fell down', but that it doesn't mean what the contextualist needs it to mean.

would dream of telling you that what he meant and was implicitly stating was that he knew that p relative to such-and such-standard” (Schiffer 1996: 326-7).

As to method, when stating claims about the truth values of [knowledge] attributions made in various contexts, contextualists rely heavily on intuitions, mainly their own. Although I won't be stressing this methodological question, it is worth asking how reliable and robust such intuitions are, why we should assume that they are representative of people's intuitions in general, and why we should take them to provide evidence about the meaning of 'know' and the semantic contents of knowledge-ascribing sentences. Nichols, Stich, and Weinberg (2003), after making a series of empirical studies of people's intuitions about various epistemologists' examples, conclude that epistemic intuitions are not nearly as universal or robust as contextualists dogmatically assume. Our own experience tells us similar things. For instance, we all know people who insist that they “knew” things that they now acknowledge to be false. So does knowledge not even entail truth? There are college administrators who describe universities as repositories and transmitters of “knowledge,” regardless of how much of what passes for knowledge is true (or adequately justified, for that matter). There are cognitive psychologists concerned with the “representation of knowledge,” whether or not what is thus represented is true. And there are sociologists (of knowledge) who study how “knowledge” (true or not) is distributed and manipulated, and many of them don't even think there is such a thing as truth. Now contextualists, like other epistemologists, would balk at these uses of 'knowledge'. They would insist that administrators, psychologists, and sociologists use the term loosely, as if it meant 'what passes for knowledge', which it doesn't. In so doing, they would be debunking the semantic intuitions of all those who use the term 'knowledge' in this allegedly loose way. I would agree with them. But on what grounds can they, as *contextualists*, dismiss these intuitions? How, on contextualist grounds, are they to decide which intuitions to rely on and which to debunk?¹⁸

Also, it is worth keeping in mind that most of the time, outside of epistemology, when we consider whether somebody knows something, we are mainly interested in whether

¹⁸ I am not suggesting that there is no basis, though in my view people's seemingly semantic intuitions are neither reliable nor robust. For one thing, they can be insensitive to the difference between the semantic content of an uttered sentence and what is implicit in the speaker's uttering of it (Bach 2002) or even what the speaker implicates (Nicholle and Clark 1999).

the person has the information, not in whether the person's belief rises to the level of knowledge. Ordinarily we do not already assume that they have a true belief and proceed to focus on whether or not their epistemic position suffices for knowing. Similarly, when we say that someone does not know something, typically we mean that they don't have the information. So the examples contextualists use to make their case, to drive their intuitions and ours, are not representative of the knowledge attributions that people ordinarily make and the concerns people have in making them.

I will not dwell on the questions raised for contextualism in the last few paragraphs. Leaving aside the linguistic and methodological difficulties for contextualists to overcome, the real question is whether they have provided reason to suppose that there are many [knowledge] relations, each involving a different epistemic standard. So far we have seen that even if contextualism is correct, so that which proposition a simple 'knows'-ascription expresses in a given context is determined by the operative epistemic standard, this does not mean that the other propositions it can express in other contexts somehow go away. They can be expressed in any context by more explicit knowledge-ascribing sentences, in which 'knows' is explicitly indexed or relativized. This point will be relevant to assessing the contextualist strategy for resolving skeptical paradoxes.

II The Contextualist Strategy

Contextualists suppose that the epistemic standard operative in a given context affects people's intuitions regarding the truth or falsity of a simple 'knows'-ascription as uttered in that context. They think this alleged empirical fact can be explained by the semantic fact (if it is a fact) that a given 'knows'-ascription can express different propositions in different contexts. It would help explain the psychological fact (if it is a fact) that different propositions expressible by the same sentence come to mind in different contexts. Of course the truth value of these propositions, each of which is expressible (by an elaborated 'knows'-ascription) in any context, is another matter. In this section we will consider how contextualists deploy their thesis to neutralize skeptical arguments.

Contextualists try to resolve skeptical paradoxes by reconciling the immovability of common sense with the irresistibility of skeptical arguments. Part of their strategy is to explain why these arguments are so seductive. However, their aim is not to refute such

arguments but merely to contain them. Different contextualists consider slightly different skeptical arguments, but let's focus on just one of them. It is as representative as any. Suppose we make the naive statement that a certain George knows that he has hands. Neither he nor we have considered the possibility that he's a BBIV, a bodiless brain in a vat (one with a body might have hands), but then a skeptic presents us with an argument:

Skeptical Argument

If George knows that he has hands, then George knows that he isn't a BBIV.

George doesn't know that he isn't a BBIV.

George doesn't know that he has hands.

Contextualists don't rebut the argument directly, by denying its validity or rejecting a premise. Their strategy is more subtle, to expose a sneaky kind of equivocation. The equivocation is not within one Skeptical Argument but across arguments. That is, the form of what appears to be a single Skeptical Argument masks a multitude of distinct arguments. These arguments are all valid, the contextualist grants, but none of them has drastic skeptical consequences. In most cases the argument is unsound; it is sound only in the extreme case, but there it is of little consequence. Specifically, the sentence comprising the second premise expresses different propositions in different contexts, and it is false in most of them.¹⁹ It is true only in what I'll call a *skeptismic* context, where skeptical standards prevail.²⁰ So the contextualist concedes that the argument is sound, but only in a skeptismic context, where far-fetched possibilities run rampant, possibilities that ordinarily may be ignored.²¹

¹⁹ Contextualists generally agree that the first premise, though it too expresses different propositions in different contexts, is true in all contexts. Heller (1999b) is an exception — he rejects relativized closure. However, it should be noted that although the Skeptical Argument is generally assumed to rely on closure, that is not quite accurate. For one could defend the first premise not by applying a closure principle but by arguing that knowledge requires that one's evidence eliminate all alternatives. Also, as Harman and Sherman (forthcoming) have argued, the intuitions that seem to support closure really support only the weaker claim that knowing requires justifiably and truly taking for granted that no counterpossibilities obtain.

²⁰ I use the neologism 'skeptismic' rather than 'skeptical' to avoid any appearance of endorsing the Skeptical Argument, even in respect to a so-called skeptical context. It would be inaccurate to call them 'epistemological' contexts, since there are plenty of epistemological contexts that don't concern skepticism.

²¹ Whether a skeptical argument actually creates a skeptismic context is another matter. As DeRose points out, "a contextualist can provisionally assume a skeptic-friendly version of contextualism, leaving it an open question whether and under which conditions the skeptic actually succeeds at raising the standards"

OK, we make an ordinary statement in an ordinary context by saying, “George knows that he has hands” (actually, this common example is a bit far-fetched, since it is not the sort of statement we would ordinarily make — almost everybody who has hands knows that). Then a skeptic confronts us with the Skeptical Argument. Contextualists contend that as soon as he does that, he has sneaked in a change of context. Since the first premise is true in both ordinary and skeptic contexts (not that its content is the same in both), this happens when he asserts the second premise. So, by the time we get to the conclusion, the skeptic has presented us with a compelling argument, indeed a sound one *in that context*. But we don’t realize that he has shifted the context on us. So we don’t realize that what he has argued for does not conflict with what we initially asserted. Indeed, the skeptic does not realize this either, since he thinks that he has refuted what we said, not changed the subject. It is only after we (and he) receive the contextualist revelation that we can appreciate that a change of context has occurred. At that point we are no longer seduced by the Skeptical Argument: we can concede its soundness in skeptic contexts without losing confidence in the [knowledge] attributions we make in ordinary contexts.

It is easy for contextualists to misrepresent what they are claiming about the Skeptical Argument. For example, look at how David Lewis describes the situation:

When we do epistemology, we make knowledge vanish. First we do know, then we do not. But I had been doing epistemology when I said that. The uneliminated possibilities were not being ignored — not just then. So by what right did I say even that we used to know? In trying to thread a course between the rock of fallibilism and the whirlpool of scepticism, it may well seem as if I have fallen victim to both at once. For do I not say that there are all those uneliminated possibilities of error? Yet do I not claim that we know a lot? Yet do I not claim that knowledge is, by definition, infallible knowledge? I did claim all three things. But not all at once! (1996: 566)

Here and throughout his paper, except at the very end (see note 4 above), Lewis commits some intentional use-mention conflation (“to get my message across I bent the rules”). Semantic ascent would have prevented that, but then he would have not been able to get his message across. In any case, knowledge doesn’t vanish on account of epistemology.

(1995: 6). This does not question the soundness of the skeptic’s argument *if* the skeptic succeeds at raising the standards.

As Mark Heller clearly explains, when uneliminated possibilities are brought up and the standards are raised, “It is misleading to describe this as a loss of knowledge. Even after the skeptic changes the standards on us, S still has the property that she had before the change of standards. There is no property that she loses” (1999a: 121). Certain knowledge-ascribing sentences go from being true to being false, but only because they express different propositions from one context to another. For Lewis this a matter of which possibilities are “properly ignored,” and that can vary with the context.

DeRose recognizes that the contextualist account of how this can be “involves the standards for knowledge being changed in a conversation” (1995: 6). So he rightly raises the question of why the Skeptical Argument “can be so appealing when considering it in solitude, with nothing being said.” In this situation there is no one else to raise the standards, and no context other than the context of one’s thinking. Even so, DeRose suggests that “there is a rule for the changing of the standards for knowledge that governs the truth conditions of our thoughts regarding what is and is not known that mirrors the [one] for what is said” (1995: 7). It is hard to see how this could be so, for in one’s thinking one could perfectly well entertain thoughts that explicitly represent the strength of standard that indexes or relativizes ‘know’. One could explicitly think thoughts with the contents of ordinary [knowledge] attributions or, just as easily, explicitly think ones with the contents of skeptisemic [knowledge] attributions. One’s context does not prevent one from doing both. Of course, contextualists are not suggesting that ordinary folk are contextualists. So it wouldn’t occur to people to think these things explicitly. Even so, the relevant thoughts people can think are explicitly expressible by means of elaborated (indexed or relativized) ‘knows’-ascriptions.

So the contextualist diagnosis of how skeptical arguments fool us does not apply when these arguments are framed in terms of elaborated (indexed or relativized) ‘knows’-ascriptions. To be deceptive, these arguments have to involve simple ‘knows’-ascriptions, such as ‘George knows/doesn’t know that he has hands’. Only then could it be easy, due to an implicit shift in standards (on the contextualist diagnosis), to conflate the contents of different attributions made with the exact same words.

Contextualism is clearly an error theory. As Stephen Schiffer explains, skeptical puzzles arise because “people uttering certain knowledge sentences in certain contexts

systematically confound the propositions that their utterances express with the propositions that they would express by uttering those sentences in certain other contexts” (1996: 325). Schiffer finds this implausible (whether the claim is that ‘know’ is ambiguous, indexical, relative, or vague) because, for example, a Moorean and a skeptic can understand each other’s utterances (and indeed their own utterances). So they should be able to recognize any shift in the content of the same sentence (or its negation) as uttered before and after the change in standard.²² But, according to contextualism, they don’t, at least not prior to hearing about contextualism. For example, if a Moorean dogmatically utters ‘George knows that he has hands’ and a skeptic springs the Skeptical Argument on him, the Moorean doesn’t recognize that the skeptic isn’t really contradicting him, and the skeptic doesn’t either. Neither recognizes that the skeptic has changed the subject. Not only does the Moorean not realize he’s being duped, the skeptic doesn’t realize he’s duping him.²³

The contextualist story is that people get fooled because they don’t notice when the bar gets raised. However, as we saw in the previous section, we ought to be able to make explicit what the different propositions are which, according to contextualism, can get expressed by the same simple ‘knows’-ascription as used in different contexts. And once we do that, there is nothing to get fooled about. As we will see next, there is more to skeptical arguments than meets the contextualist’s eye.

III Contextualism and Skepticism

²² The situation would be like what happens when someone in one time zone asks or tells another what time it is. One could imagine a similar conversation about weight between an earthbound person and a man on the moon. For discussion of Schiffer’s objection and how the contextualist might reply, see Hofweber 1999.

²³ In response to Schiffer’s argument, Cohen (2001) contends that contextualism is an error theory only with regard to “meta-judgments” that different utterances of the same [knowledge] attributing sentence have the same contents. But surely, if people fail to recognize a shift in content between two utterances of the same sentence, or mistakenly detect a contradiction when ‘not’ is included in one, they’ve got the content one of the utterances wrong. For example, the Moorean either misunderstands what the skeptic says or misunderstands what he himself said. Ram Neta (2003), who recognizes that Cohen’s attempt to kick Schiffer’s objection upstairs is unsuccessful, urges the contextualist “to develop a version of contextualism that helps us to appreciate the semantically relevant difference between the context in which Moorean anti-skepticism is false and the context in which it is true, and thereby frees us from puzzlement.” Our devices of explicit indexing and explicit relativization do just that.

Contextualists aim to diagnose and relieve the intuitive tension generated by the clash between the deliverances of common sense and the seductiveness of skeptical arguments. As Cohen makes clear, contextualists do not intend their efforts at resolving skeptical *paradoxes* to be taken as refutations of skeptical *arguments* (1999: 69). DeRose acknowledges that “in claiming that my belief that I have hands is sensitive, I betray my conviction that I am not a BIV in this world or in nearby worlds” (1995: 50). So there is no point in accusing contextualists of begging the question against skepticism.²⁴ Still, many philosophers have complained that contextualists do not really come to grips with the force and content of skeptical arguments (see Feldman 1999 and 2001, Klein 2000, Kornblith 2000, and Sosa 2000). The complaint is simple: the contextualist’s attempt to marginalize skeptical arguments by restricting them to skeptical contexts ignores the fact that skepticism denies that we have knowledge even by *ordinary* standards. As Richard Feldman writes,

The question skepticism raises is about whether our evidence really is good enough to satisfy the standards for knowledge. One can think that the familiar skeptical possibilities introduce grounds for doubt that defeat our evidence for our ordinary beliefs. One can think that we have no evidence at all that favors our ordinary beliefs rather than their skeptical rivals. Either way, there’s reason to wonder whether we really do satisfy the ordinary standards. The debate about skepticism is thus seen not as a debate in which the quality of our evidence is agreed to and the debate results from differing views about what the standards for knowledge are. Instead, it is a debate about how good our evidence is. Understood that way, it’s difficult to see the epistemological significance of decisions about which standards are associated with the word “knows” in any particular context. (unpublished reply to DeRose forthcoming)

Feldman’s point, then, is that when a skeptic brings up far-fetched possibilities and argues that we can’t rule them out, he is not raising the standards for what it takes to belong to the extension of the word ‘knowledge’. Rather, he is using these possibilities to show that it is much tougher than we realize for a belief to qualify as knowledge at all, even by the normal standards governing ordinary contexts, that is, to have the property that the word ‘knowledge’ actually and ordinarily expresses. So contextualists haven’t really addressed what Kornblith calls “full-blooded” skepticism.

²⁴ Contextualists think that because, as Heller proclaims, “it is a completely convincing response to the skeptic” to point out that “even after the skeptic changes the standards on us, [the subject] still has the property that she had before the change of standards” (1999a: 121).

Skeptics are not proposing to reform the meaning of the term ‘knowledge’. They are recommending that we use it to mean what it ordinarily means but use it much more carefully, even if it turns out rarely to apply. So it is wrong to charge them with proposing that the word ‘knowledge’ be used so that it expresses a property that is much more difficult for a belief to have than the property that term actually expresses. This would be the tired complaint of an ordinary-language philosopher who, preferring to speak with the vulgar, would accuse the skeptic of inventing a special, philosophical sense of ‘know’. Rather, the skeptic is suggesting that we make a much more serious effort than we ordinarily do at ascribing that property accurately and let the chips fall where they may.

Just imagine that there were lots of fool’s gold around and that people commonly described it as ‘gold’ simply because it looked and felt like gold. An auric skeptic, someone who doubted that most of the stuff that passes for gold really is, might advise us to be much more careful about what we count as gold (“All that glitters is not gold”). In a way, he would be suggesting a reform in our use of the term ‘gold’, but this would concern how we go about applying the term, not what the term actually applies to. He would be advising us to make sure to apply it only to samples of gold. The skeptic is offering similar advice about using the word ‘knowledge’. He is not recommending that we use it to express a more precious property but that we be much more careful about what we take to possess the precious property that ‘knowledge’ actually expresses. He is concerned with knowledge, not ‘knowledge’.

Now there is a relevant difference between the auric skeptic and the epistemic skeptic. In the above scenario the auric skeptic would not be unduly demanding — he would be right. However, the epistemic skeptic is wrong, or so we non-skeptics believe. He is wrong to suppose that knowledge is limited to beliefs that are justified to the highest degree, beliefs the evidence for which eliminates all counterpossibilities, however far-fetched. That is why we think the epistemic skeptic is being too demanding. Even so, as he see things, our situation is analogous to that of the auric skeptic in the above scenario. To appreciate this, change that scenario and suppose that we are in fact very good at identifying samples of gold as gold and at not mistaking fool’s gold for the real thing. And suppose there’s very little fool’s gold around. Our auric skeptic would still

insist that we be much more careful about what we deem to be gold, but in this case he would be too demanding. He could raise the possibility of widespread fool's gold, but in so doing he would not be alerting us to a real possibility. In letting his fantasies rip, he would be much like the epistemic skeptic.²⁵

Contextualists grant that we can't rule out, and that our evidence doesn't eliminate, global skeptical possibilities (dreaming, demon, BBIV, Matrix, etc.), and they concede that this makes it impossible or at least very difficult for us to [know], relative to skeptical contexts, what ordinarily we take to be everyday empirical facts. Nevertheless, they have no qualms about assuming that people easily [know], relative to ordinary contexts, those very same facts — even though our evidence regarding skeptical possibilities is the same.²⁶ It's just that ordinary standards don't require that we rule out or that our evidence eliminate such "irrelevant" possibilities. Relative to those standards we can [know] that such possibilities don't obtain even though our evidence doesn't rule them out. However, no self-respecting skeptic is going to concede that there is any context relative to which ordinary [knowledge] attributions can be true even though the subject's evidence does not rule them out. Indeed, he has no reason to concede that it rules out any counterpossibilities if it doesn't rule out the skeptical ones.

What would the skeptic be willing to concede regarding knowledge attributions made when less demanding standards are operative? He might concede, hypothetically, that *if* someone knew that no skeptical possibilities obtain, he *would* know ordinary empirical propositions to be true.²⁷ Since we ordinarily presuppose these things, we think that the knowledge attributions we make according to less demanding standards are often true. However, that does not make them true, even relative to those standards. The skeptic does not buy the contextualist's claim that [knowledge] attributions are true by ordinary

²⁵ In this paragraph I am ignoring another relevant difference between the auric skeptic and the epistemic skeptic. The epistemic skeptic's general principles always take precedence over paradigm applications of the word 'knowledge', but not because it is a clear case of a natural kind term. Whereas the auric skeptic takes for granted that there is an objective, uncontroversial standard for gold, there is no such standard for knowledge, even assuming that invariantism is correct.

²⁶ Heller (1999b) is the one exception I know of. He does have qualms — he rejects Closure.

²⁷ As Jonathan Schaffer has pointed out to me, noting that contrapossible counterfactuals are hard to evaluate, this would be an odd concession for the skeptic to make, since he denies that there is any possible world in which one could know, or even or correctly and justifiably take for granted (see note 19), that no skeptical possibility obtains.

standards. The mere fact that people's willingness to make and accept knowledge attributions is governed by their application of ordinary epistemic standards does not mean that these knowledge attributions are generally true.

Not only do contextualists not acknowledge the full force of skepticism, they also concede too much to skepticism when they suggest that skeptical arguments are seductive because of a subtle shift to a very high-standards context. Here's an alternative explanation, from the standpoint of moderate (nonskeptical) invariantism (to be defended in section V), of why skeptical arguments are seductive. When we are presented with a skeptical argument and confronted with what we ordinarily take to be far-fetched sources of error, in effect we are asked to imagine ourselves, with our current experiences, (apparent) memories, and beliefs, plunked into a world of we know not what sort. It could be a dream world or a demon world, a BBIV or a Matrix world, or any of a whole host of others. Or it could be a world of just the sort we think we're in. But we're not supposed to have any prejudices about which sort of world we're being plunked into. Since each of the possible worlds is consistent with our having the perceptual and memory experiences and beliefs we have, there is nothing to make the world as we commonly conceive of it epistemically special in any way. It's just one of those countless sorts of worlds any one of which we could be plunked into. So of course we can't tell that we're in any one of the them, in particular, the world as we commonly conceive of it.

This explains why skeptical arguments, as inspired by Descartes's systematic doubt, are so seductive, but it doesn't show that they are any good. Yes, it's true that if we were suddenly plunked into a world, we wouldn't be able to tell what sort of world we were in. But that's not our situation. To know in this world, it is not necessary to be able to discriminate between the different possible worlds we might be in. It is not necessary to know that we're not in a world where we would be chronically prone to uncorrectable and undetectable error, at least not if knowing this requires going out and verifying that we're not in such a world (it would be impossible to verify *that*). True, a skeptical scenario would seem no less absurd if one were in it than it does in fact, but that doesn't show that it is not in fact absurd. The fact that there are possible worlds in which we would know very little does not show, or even suggest, that we are in such a world. Knowledge may not be as easy to come by as people casually suppose, but to be in a

world which is stable in various fundamental respects, with which we informationally interact in clearly explicable ways and in which we communicatively interact to transmit information successfully, is to be in a world in which there is plenty of knowledge to be had.

I don't know how well this rebuts skeptical arguments (for example, it dogmatically rejects internalism about knowledge and epistemic justification), but at least it doesn't change the subject by accusing the skeptic of changing the subject.

IV Contextualism and Skeptical Invariantism

Not surprisingly, contextualists try to bolster their case by arguing against invariantism. Unfortunately, they tend to limit their attention to *skeptical* invariantism. No doubt influenced by how Unger introduced the contrast between contextualism and invariantism (1984: 6-11), they follow him in likening the invariantist view of 'knows' (the view of Unger 1971 and 1975) to his invariantist view of 'flat'. On that view, if something is flat, nothing can be flatter than it. So the only flat things there are are absolutely flat, as flat as flat can be. Anything less flat than that can only be somewhat flat, relatively flat (as compared with something else), rather flat, or even very flat, but not really flat. It must be perfectly flat for that. So when we describe such a thing as flat, we are not speaking truly but only truly enough for practical purposes. Similarly, on the skeptical invariantist view of 'knows', you can know that p only if you are as well-positioned as possible about p as you could be about anything, e.g., that you exist. Given such a demanding view, empirical knowledge is very hard to come by.

In challenging this very demanding view, DeRose points out that "Unger did admit that varying standards for knowledge govern our use of sentences of the form 'S knows at t that p', but did not endorse contextualism because [he] claimed that these varying standards were only standards for whether it was *appropriate to say* that S knows" (1999: 192). DeRose describes this as a "warranted assertibility maneuver," or WAM.²⁸ Although he recognizes that WAMs can be legitimate, he rejects their application to knowledge attributions (1999: 196-203, and 2002: 191-94). He thinks these are as "lame"

²⁸ DeRose does not use 'warranted assertibility' to mean what it usually means in philosophy, especially in discussions of anti-realism, namely a kind of metaphysical or epistemological surrogate for truth.

as the claim that ‘bachelor’ means man and that saying that someone is a bachelor merely implicates that he’s married. DeRose gives various plausible reasons for doubting the skeptical invariantist contention that ordinary knowledge attributions are generally not true but merely appropriate to make and in that sense warrantably assertible.

DeRose is not entirely fair to Unger, who was offering not a WAM but a kind of error theory about how people use knowledge-ascribing sentences. Consider that there are two rather different ways in which a sentence can be warrantably assertible without being true: (a) what is said, that S knows that p, is not true but, because it is close enough to being true, can be warrantably asserted anyway, or (b) in saying falsely that S knows that p, the speaker implicates something else which is true. Unfortunately, DeRose does not explicitly distinguish the two. This is unfortunate because Unger, at least if he meant anything of the sort, meant (a), and what DeRose argues against is (b). DeRose finds it highly implausible that when people make ordinary knowledge attributions and say such things as that George knows that he has hands, they are really implicating or otherwise pragmatically conveying something else instead. He goes to great lengths to show how implausible this is, but this is not Unger’s view. Unger held that we often speak loosely and casually, as when we use an “absolute” term like ‘flat’ or ‘empty’ and describe a surface as flat or a container as empty. This does not imply that in so speaking, we mean something else that we are not fully spelling out. Rather, we mean just what we are saying but are construing it loosely. The flatness and emptiness attributions we ordinarily make pass as true but are not really true. To borrow a favorite contextualist phrase (but to use it differently), they “count as true.” They’re true enough for practical purposes, and that’s good enough for us.

However, it is not clear that Unger thought of knowledge attributions as even warrantably assertible in this way (in sense (a) above). Unger was not using a WAM to explain why people make literally false knowledge attributions because only the weakest sort of epistemic skepticism would concede that they are very close to being true. And he certainly was not claiming that they are warrantably assertible in sense (b). He was not suggesting that attributors, in saying something literally false, implicate or otherwise convey something true. After all, ordinary folk are not privy to skeptical arguments and do not make their knowledge attributions in defiance of such arguments. Their loose use

of 'know' is hardly self-conscious or sophisticated enough for them to intend their simple, unqualified knowledge attributions to be taken loosely or for them to recognize knowledge attributions made by others as intended to be taken loosely. Unger's explanation of why people make simple unqualified knowledge attributions is nothing like a Gricean or pragmatic account of how people can say one thing and mean something else instead.

Recently DeRose (2002) has offered a new argument for contextualism, which is notable for its unwitting use of a different notion of warranted assertibility. This notion is more epistemological than pragmatic. Here he does not mean that a sentence, even if false, is warrantably assertible if uttering it implicates something true and nothing false (this is a kind of conversational appropriateness). He means something quite different, as is clear when he sums up his new argument:

The knowledge account of assertion provides a powerful argument for contextualism: If the standards for when one is in a position to warrantably assert that P are the same as those that comprise a truth-condition for 'I know that P', then if the former vary with context, so do the latter. In short: The knowledge account of assertion together with the context-sensitivity of assertability yields contextualism about knowledge. (2002: 171)

Here he must mean that what makes 'P' warrantably assertible is that one knows that P, since that is the truth-condition for one's utterance of 'I know that P'. And, needless to say, 'P' can be warrantably assertible in this sense without being so in the other (and conversely). Also, notice that this "powerful argument for contextualism" applies only to first-person cases, in which attributor and subject are the same (almost all of DeRose's examples are of this sort). Obviously, however, what makes 'P' warrantably assertible by the attributor, that the attributor knows that P, is not the truth condition for his utterance of 'S know that P', where S is somebody else.

In this section, we have seen that, contrary to what DeRose contends, the skeptical invariantist, as exemplified by Unger, does not, and need not, resort to a WAM to make his case. However, I am not suggesting that his case is a good one. Skeptical invariantism relies on the Skeptical Argument, but that is an epistemological argument, not a semantic one. If that argument is sound, then our ordinary knowledge claims are false. But this is more skepticism about knowledge than a form of invariantism about 'knows'. To claim

that the extension of the word ‘know’ is very small is just a roundabout metalinguistic way of claiming that knowledge is very hard to come by. Moderate invariantism, at least as I will defend it, does assume that skepticism is false, but it does not rely on any substantive epistemological arguments to make its claim that the semantics of ‘know’ is invariant.

V The Obvious Alternative to Contextualism: Moderate Invariantism

Non-skeptical or moderate invariantism is defined by DeRose as “invariantism that keeps the standards governing the truth-conditions of knowledge attributions constant, but meetably low” (1999: 192). This leaves open just how demanding these standards are, and for present purposes I will keep it open. I’ll just assume that whether or not most of our ordinary knowledge attributions are true, a good many of them are, far more than even the weakest form of epistemic skepticism could allow. This leaves open that plenty of them are false by almost any philosopher’s standards and that plenty of them are debatable. No doubt we often credit people with knowledge when we shouldn’t, but then we speak casually about a lot of things (Bach 2001). Still, I will assume that much of what passes for knowledge really is, including perceptual knowledge, such as that you are sitting and that there are lines of print in front of you, and knowledge about simple facts, e.g., that California is a state and that chickens lay eggs.

DeRose assumes that the moderate invariantist, like the skeptical one, has to “chalk it all up to pragmatics” (2002, 194) and rely on a WAM to make his case. However, whereas the skeptical invariantist claims that with an ordinary knowledge attribution “we mistake the warranted assertibility of the claim for truth,” for the moderate invariantist the situation is reversed: “it’s the denial of knowledge in the high standards case that’s false but appropriate: Due to the high standards for the warranted assertibility of knowledge in place there, a positive claim that the subject knows would be unwarranted (though true), and it’s the denial of knowledge that is appropriate (though false)”²⁹ (2002: 171). This applies not just to skeptical cases but to more down-to-earth yet high-standards versions of more pedestrian cases. Compare this with what contextualism says about

²⁹ Unfortunately (see note 5), DeRose’s own discussion here is marred by the fact that he focuses entirely on first-person [knowledge] denials, as when he stresses the “glaring difference between ‘I don’t know P’ (which often becomes assertable when the standards go up) and ‘not-P’ (which doesn’t)” (2002: 191).

these cases, such as the high-standards versions of the Bank and the Airport cases. Even though the subject's epistemic position is the same, according to DeRose and Cohen when the stakes go up, the standards go up, and when the standards go up, [knowledge] attributions that would normally be true are not true, and the corresponding [knowledge] denials *are* true.³⁰ The moderate invariantist has to say that when the stakes go up, the [knowledge] attributions are still true and the corresponding [knowledge] denials are still false. But he has to say more than that, as DeRose rightly insists, for there is a “glaring difference between ‘I don’t know P’ (which often becomes assertable when the standards go up) and ‘not-P’ (which doesn’t)” (2002: 191).

As we will see, the contextualist account of such examples, cases in which standards on [knowledge] attributions are allegedly raised, ignores two complementary facts: that attributing to someone knowledge that p involves (confidently) believing that p yourself; and that denying knowledge of someone who has the same evidence you have involves being at least somewhat doubtful about p.³¹ So I would not accept DeRose's stipulation regarding the high-standards version of his Bank case, according to which the attributor denies knowledge while “remaining as confident as [he] was before that the bank will be open tomorrow” and yet concedes that he'd “better go in and make sure” (1992, 913). It seems to me that unless he's trying to placate his wife, his belief would have to be shaken at least somewhat.

Consider the contextualist characterization of the high-standards Airport case, in which Mary is unwilling to assert that Smith [knows] that the plane will stop in Chicago. Given how important this question is to Mary, not only does she refrain from attributing [knowledge] to Smith but, unwilling to take Smith's word as based on his itinerary, she goes so far as to deny that he [knows] that the plane will stop in Chicago. According to

³⁰ For the sake of discussion, we're assuming that in the low-standards versions of these cases, the relevant attributions are true. You don't have to be a skeptic to think that these standards are too low and that the subject does not [know] even without the standards being raised. If you think that, then pick a different example, such as one involving ordinary perceptual identification or recollection of a simple fact.

³¹ It is interesting to note that when introducing skeptical invariantism, Unger (1971) focused not on the strength of the subject's epistemic position but on the strength of the subject's belief. He did not stress the ultra-high standards which, according to contextualists, the Skeptical Argument purports to demand. Rather, he stressed the strength of the doxastic condition on knowledge, arguing that it requires “absence of doubt or doubtfulness.” One could take a less demanding position but still insist that the doxastic condition on knowledge requires more than mere belief. There is also the question of how much conviction belief itself requires.

the contextualist, that's because it isn't true that he [knows] this. But the moderate invariantist has to say that if Smith knows in the normal, low-standards case, he knows in the high-standards case too, even if Mary is not prepared to say that he does. So, does the moderate invariantist need to rely on a WAM to explain this? Should he argue that since Smith is well enough positioned for an utterance in an ordinary context of "Smith knows that the plane will stop in Chicago" to be true, then when Mary asserts its negation in a context of heightened interest, she must be saying something false and pragmatically conveying something true?

No! Mary is making a mistake, albeit a very understandable one. Mary does not say "Smith knows that the plane will stop in Chicago" and goes so far as to assert its negation because of her own doxastic situation. Because she is not sure Smith's itinerary is reliable, she herself is not confident enough to believe that the plane will stop in Chicago. But since she is not confident enough to believe that, she can't coherently attribute knowledge of it to Smith, not if knowledge implies truth. In general, you can't coherently assert that someone else knows that p if you are not confident that p and think that it still needs to be verified. That is why Mary can't very well assert that Smith knows that the plane will stop in Chicago. Not only that, she has to deny that she knows it, since she thinks it is not yet established. And, since Smith has no evidence that she doesn't have, she must deny that he knows it either.

What is decisive here is not the attributor's lack of belief but her raised threshold for (confidently) believing.³² By this I mean that before believing the proposition in question, at least with the confidence and freedom from doubt necessary for knowing (see note 31), the attributor demands more evidence than knowledge requires. So, in the high-standards version of the Airport case, what happens when Mary double- and triple-checks and confirms to her satisfaction (in the Airport case) that the plane will stop in Chicago? She will then believe that it will stop there and will think that she knows this. However, she still won't concede that Smith knows this and indeed will still deny that he does, given that his epistemic position is no better than hers was. Now the explanation for her denial is not that she doesn't believe it herself but, rather, that her threshold of confidence has gone up.

³² Thanks to Jessica Brown, John MacFarlane, and Jonathan Schaffer for urging me to spell this out.

One's threshold for (confidently) believing a proposition is a matter of what one implicitly takes to be sufficient reason to believe it (again, if believing is compatible with residual doubtfulness, I mean the sort of confident belief required for knowing). I say "implicitly" because people generally do not reflect on such things. Even if in fact one is in a position to know something, thinking one is not a position to know it is enough to keep one from believing it (at least not without reservations) and to lead one, if it matters enough, to look into it further. When one does look further and verifies the proposition to one's satisfaction, one implicitly takes oneself now to be in a position to know it and continues to regard one's prior, weaker position as inadequate. So one cannot consistently take someone else, who was in and still is in that weaker position, to know it. In consistency, one must regard him as not knowing it.

It might seem that I have merely described in different terms what the contextualist describes as raising the standards on a [knowledge] attribution. However, what I have described is what it takes for an attributor coherently to make a knowledge attribution to someone else who has certain evidence, given the attributor's doxastic stance relative to the same evidence. Here's a way to put the difference between the contextualist view and my own. Remember that attributing knowledge that *p* requires believing that *p*. So, I am suggesting, willingness to attribute knowledge does not track the standards on the truth of a [knowledge] attribution; rather, it tracks one's threshold of doxastic confidence. In the so-called high-standards cases, the attributor's doxastic threshold goes up to the point that without additional evidence she implicitly, but mistakenly, thinks she is not a position to know. This makes my account a kind of error theory, but only minimally. It has people sometimes denying knowledge of people who have it, but it does not have them generally confused about knowledge or 'knowledge'. In high-standards Airport and Bank cases, a special practical interest gives the attributor reservations about the truth of the proposition in question and, accordingly, raises her bar for attributing knowledge to someone else. Even so, the subject knows.

Now let us consider what happens when a skeptical possibility is raised. It could be a general skeptical possibility, such as victimization by an Evil Demon, or one specific to the case, say an imagined rumor that disgruntled travel agents are distributing inaccurate itineraries. However, merely raising such a possibility, without making it plausible, does

not turn a true knowledge attribution into a false one. Making it salient is not enough. That, as Patrick Rysiew (2001) has shown, affects at most the assertibility of the knowledge attribution, because making the attribution pragmatically conveys that a newly raised possibility has been ruled out.³³ Raising plausible possibilities, on the other hand, indicates real doubts on the part of the attributor and, if taken seriously, lowers the doxastic state of his audience. Moreover, if these plausible possibilities are objective possibilities, ones that bear on the subject's epistemic position, and if the subject's epistemic position is not strong enough to rule them out, then the subject does not know that *p*, quite independently of the attributor's context. In no case, then, is the truth condition or truth value of a knowledge attribution affected by the epistemic standards that prevail in the context of attribution. All that is affected is the attributor's willingness to make it and the audience's willingness to accept it, by way of raising their threshold for (confidently) believing.

We are now in a good epistemic position to reply to DeRose's contention that moderate as well as skeptical invariantism requires the use of WAMs in order to explain intuitions about ordinary and high standards [knowledge] attributions. He writes,

Invariantists do not begin with a good candidate for WAMing, and they have to explain away as misleading intuitions of truth as well as intuitions of falsehood. For in the "low standards" contexts, it seems appropriate and it seems true to say that certain subjects know and it would seem wrong and false to deny that they know, while in the "high standards" context [as in the Bank case], it seems appropriate and true to say that similarly situated subjects don't know and it seems inappropriate and false to say they do know. Thus, whichever set of appearances the invariantist seeks to discredit — whether she says we're mistaken about the "high" or the "low" contexts — she'll have to explain away both an appearance of falsity and (much more problematically) an appearance of truth. (DeRose 2002: 193)

³³ As Jessica Brown (forthcoming) points out, before drawing conclusions from their examples contextualists need to control separately for salience and for practical interest. She argues that salience alone does not raise the standards, at least not in the clear way that practical interest does, and she uses this observation to develop a non-skeptical version of invariantism. Her version is a modification of Patrick Rysiew's (2001), who thinks it is salience which affects the knowledge attributions people are willing to make and which bears not on the truth or falsity of the attributions but on what they pragmatically convey. Brown, like Rysiew, employs sophisticated WAMs, not lame ones of the sort that DeRose thinks moderate invariantism is stuck with. My version of moderate invariantism does not rely on WAMs at all.

The problem is that DeRose accepts the appearances at face value.³⁴ Moderate invariantists should accept intuitions about ordinary [knowledge] attributions at face value but should reject DeRose's intuitions about the "high standards" Bank case (where the cost is high of the bank not being open on Saturday). The attributor's high stakes (on Friday) when asserting or accepting as true 'Keith knows that the bank is open on Saturday' do not translate into higher standards for its truth. Rather, she has good practical reason, because of the cost of him being wrong, not to take Keith's word for whether the bank is open on Saturday. Given that, she doesn't accept his statement as true without checking further. So she can't consistently accept or assert 'Keith knows that the bank is open on Saturday' as true.

Moderate invariantists should also reject the intuition that [knowledge] denials involving skeptical possibilities are true. To repeat the anti-internalist sentiments expressed at the end of section III, I think the moderate invariantist should not concede that there is something right about the intuition that George does not know he is not a BBIV and that an utterance of 'George does not know that he is not a BBIV' is true, at least in a skeptistic context. Rather, he should insist that George does know he is not a BBIV and that the intuition that he doesn't is based on the false assumption that in a skeptical scenario George's epistemic situation would be no different. To be sure (*pace* Williamson 2000), George does not have evidence that he would not have if he were a BBIV, but that doesn't matter. The intuition that some people have that it does matter seems to be based on a leap, from the obvious truth that George, if he were a BBIV, wouldn't know it and would still believe that he is not a BBIV, to the conclusion that in fact he doesn't know he's not a BBIV.³⁵ If he were a BBIV, there would be lots of things he wouldn't know, even if the world were otherwise as much as possible like the actual world, and certainly if it were vastly different. His beliefs are insensitive to the difference. But he can be in a position to know things about the actual world, such as that he has hands and that he is not a BBIV, even if, were the world quite different (or if the

³⁴ As mentioned earlier, Nichols, Stich, and Weinberg (2003) have shown that epistemic intuitions are not nearly as universal or robust as contextualists dogmatically assume.

³⁵ Here I am ignoring possible content-externalist differences in his beliefs, hence in what he represents his evidence to be. It is an interesting question whether issues concerning content internalism and externalism have any bearing on the debate between epistemic internalists and externalists.

causes of his beliefs were quite different, as in a benign-demon world), he wouldn't know very much about it. Only certain sorts of worlds and relations to the world are such that one can know things about that world. The prevalence of massive error in some possible worlds, especially in worlds remote from this one, does not show the real possibility of massive error in the actual world.

Skeptical invariantism is admittedly an error theory, and contextualism is clearly an error theory too. Is my version of moderate invariantism also an error theory? It is, but only in a minimal way. According to skeptical invariantism, people commonly make false knowledge attributions. According to contextualism, people commonly fail to recognize shifts in the contents of 'knows'-ascriptions and thereby sense contradictions that are not there. It implies that people are frequently unaware of differences in the contents of [knowledge] attributions. The only sort of error that my version of moderate invariantism attributes to people, other than the error of being temporarily taken in by skeptical arguments (attributing this error is not specific to moderate invariantism), is one of excessive caution when it comes to believing things with confidence.

VI Ignoring as Evidence

There is more to be said about the interaction between one's doxastic state and one's epistemic position and about the implications of that for what one is prepared to say about someone else's epistemic position. Pretty much everyone, contextualist and invariantist alike, agrees that knowing that *p* requires that one's experience/evidence/justification rule out counterpossibilities (alternatives to *p*, threats to the basis for one's belief that *p*). There is plenty of disagreement about how best to formulate this, especially if one rejects the skeptic's contention that knowing requires ruling out all such possibilities. It is common to limit the requirement to ruling out relevant alternatives, and there are different variations on this approach. Typical internal problems for such an approach include spelling out what it is to rule out an alternative, whether it is the subject or his evidence that does this, and, of course, what it is for an alternative to be relevant. It will be instructive to focus on David Lewis's account.

Lewis exhorts you to "do some epistemology [and] let your paranoid fantasies rip!" (1996: 559). OK, let 'em rip. That's what Descartes did with his Evil Demon fantasy, and

the BBIV (or the “Matrix”) scenario is just a high-tech version of that. But just imagining yourself in such a scenario is not to take seriously the possibility that you’re in one. These so-called skeptical “hypotheses” are just fantasies. Getting yourself and your conversational partner to entertain such fantasies may change the context but it doesn’t turn them into real possibilities. It would seem, then, that they can be safely ignored. But for Lewis things are not so simple: “Our definition of knowledge requires a *sotto voce* provision. S knows that P iff S’s evidence eliminates every possibility in which not-P — Psst! — except for those possibilities that we are properly ignoring” (1996: 554).³⁶ Any possibility compatible with the experience (with its having the content that it has) is not eliminated.³⁷ But there will always be skeptical possibilities, as many and varied as you can dream up, that are compatible with your experience. So if skepticism is to be avoided, they can’t count against the truth of ordinary [knowledge] attributions. This is possible only if they can be properly ignored without having to be eliminated. However, by Lewis’s Rule of Attention, “a possibility not ignored at all is *ipso facto* not properly ignored” (1996: 559). So a skeptical possibility once presented cannot properly be ignored. Because it is not eliminated by one’s experience, according to Lewis it inevitably “destroys knowledge.” As even Lewis’s fellow contextualists acknowledge, this requirement makes it too easy for the skeptic: he can prevail just by mentioning far-fetched possibilities. But there is more to the Rule of Attention than this.

Suppose that it is not an experience but the person having it that eliminates a possibility (this will simplify our description of the situation). In that case, if the thought of some possibility occurs to you, you have to rule it out — you can’t just disregard it. On the other hand, it would seem that you can’t rule out a possibility if the thought of it doesn’t occur to you. If it doesn’t occur to you that there might be cleverly painted mules in the vicinity, you can’t very well rule that out. Now could this keep you from being in a

³⁶ Right before giving this definition, Lewis remarks that ‘an idiom of quantification, like every’, is normally restricted [semantically?] to some limited domain’ (1996: 553), so that ‘every possibility in which not-P’ does not include those that are being ignored. “They are outside the domain, they are irrelevant to the truth of what is said.” But at the very least (never mind whether this restriction is semantic — see Bach 2000), being ignored, which suffices for being outside the domain, can’t be irrelevant to the truth of what is said. These possibilities must be properly ignored, as required by the definition of knowledge.

³⁷ Here, an experience or memory that P “eliminates W iff W is a possibility in which the subject’s experience or memory has content different from P” (1996: 553). Notice that on Lewis’s conception of elimination, it is the experience, not the person having it, that eliminates a possibility.

position to know you're looking at a zebra? Ordinarily you don't have to rule out such a possibility — you have no reason to and there is nothing about your environment that requires you to — and the thought of it doesn't even occur to you. But what if the thought of such a possibility did occur to you? If it is as far-fetched as this one, and you have no reason to think it isn't, can't you just dismiss it? The mere fact that the thought of it occurs to you shouldn't make any difference. Or should it?

Offhand, it might seem that whether or not the thought of a certain possibility occurs to us has no epistemic significance and that what matters is what we do. Possibilities just occur to us, and we should take realistic ones seriously and do what it takes to rule them out. We can just dismiss the far-fetched ones if and when they occur to us. Thinking of them would be a distraction and, if chronic, a nuisance, but that would be all. In fact, however, possibilities don't just occur to us at random. Insofar as our cognitive processes work efficiently and effectively toward our cognitive goals, the fact that a possibility occurs to us provides evidence that it is worth considering. Not only that, the fact that a possibility does not occur to us provides evidence that it isn't worth considering (such evidence is highly defeasible, since it may be our ignorance that keeps the thought of a relevant possibility from occurring to us). In this way, we can safely jump to conclusions without having to verify the countless implicit assumptions that we make in our everyday reasoning.³⁸

If our cognitive processes are operating well, generally the thought of a possibility contrary to something we're inclined to take as fact occurs to us only if it is a realistic possibility, not a far-fetched one. We can always conjure up wild possibilities, as in flights of skeptical or paranoid fantasy, but when we're engaged in normal inquiry or just trying to identify what we're perceiving or remember some bit of information, we take into account only those counterpossibilities that sometimes arise in situations of the sort we're concerned with. So it is the very occurrence of the thought that gives us a reason for considering the possibility being thought of. No wonder, then, that “a possibility not ignored at all is *ipso facto* not properly ignored”! Of course, should it occur to us we may find reason to dismiss it.

³⁸ I defended this conception of default reasoning in Bach 1984 and used it to defend a form of reliabilism about justified belief in Bach 1985.

Here's a simple example of how the unbidden thought of a possibility can undermine one's knowledge by shaking one's belief. Someone once asked me to name the capital of Kentucky and I immediately thought, "I know that. It's Frankfort." But then it occurred to me that, well, maybe it's Lexington. This put doubt in my mind, and I was more inclined to say, "I'm pretty sure it's Frankfort, but it might be Lexington." If asked to choose, I would have said it's Frankfort, but the very occurrence of the thought that it might be Lexington undermined my knowledge, at least temporarily, by shaking my belief. Second thoughts rightly yield doubt, but most of the time we simply rely on the reliability of our memory (or our eyesight or whatever) and don't have second thoughts.

We rely on our reliability at thinking of counterpossibilities when they are worthy of consideration and at knowing when to look further before settling into a belief. Also, if the relevant cognitive processes are functioning well, the *nonoccurrence* of the thought of a certain counterpossibility provides evidence, albeit highly defeasible, that this counterpossibility is *not* worth considering. However, we cannot explicitly *take* the nonoccurrence of the thought as the evidence that it is, for that would entail thinking of the counterpossibility in question. We cannot explicitly weigh the evidence that the nonoccurrence of the thought provides, at least not at the time (in retrospect we may reason that if something was a realistic possibility, it would have occurred to us).

Suppose, contrary to the above picture, that possibilities contrary to something one is otherwise disposed to believe came to mind independently of any evidence one has for their actually obtaining. If our minds worked that way, then the occurrence of the thought of a counterpossibility would just be a nuisance. It would be like random, unbidden thoughts that one has left the front door unlocked or the headlights on. One caters to such thoughts and one checks, even if it is inconvenient to do so, in order to make the thought go away. So it would be with random, unwarranted thoughts of counterpossibilities to something one is otherwise disposed to believe. But this is not how our minds generally work. Thoughts of counterpossibilities occur to us generally because there is reason to consider them, hence a need to rule them out. We need not have specific, articulable evidence, but somehow, at least to some extent, our belief-forming processes are tuned into plausible sources of error. If something like this picture is correct, then the nonoccurrence of the thought of a counterpossibility is evidence that the

counterpossibility does not obtain, but it is not evidence we can directly consider. Instead, we rely on the reliability of our tendency to think of counterpossibilities when and only when they're worth considering.

Now what does this suggest about knowledge attributions and the possibilities that come to mind or get brought up in a context of attribution? Let's say that a scenario is epistemically irrelevant to a knowledge attribution if the mere possibility of its obtaining does not affect the truth of that attribution. A scenario can be epistemically irrelevant because it is just a wild skeptical fantasy, whether global or specific to the case, or because, despite the fact that it is something the attributor needs to rule out for practical reasons or for bad skeptical reasons, it has no bearing on the truth of the knowledge attribution. However, considering what is in fact an epistemically irrelevant scenario gives the attributor reservations about believing the proposition in question and puts the attributor in the position of having to deny that he knows. And, as we saw in the previous section, this is enough to put her in the position of refraining from asserting that the subject knows and even of falsely denying that she knows. What does this, we can now see, is the consideration of a possibility that epistemically is not worth considering. That is enough to keep one from settling into a belief one would ordinarily adopt. Ordinarily, when the thought of such a possibility does not occur, one would form the belief on the basis of the evidence one already has.

VII Summing Up

Contextualism is the thesis that a sentence of the simple form 'S knows at t that p' can be true as uttered in one context and false as uttered in another, depending on the epistemic standards that govern the context. The standards governing the context help determine which knowledge-attributing proposition the sentence expresses in that context. This means that there is no one knowledge relation and that the different propositions expressible by such a sentence involve different [knowledge] relations. Given this understanding of what contextualism says, I have argued for several things:

I. Even if contextualism is true, so that a simple knowledge-ascribing sentence can express different propositions in different contexts, those different propositions are not themselves context-bound. Each such proposition can be expressed by a more elaborate

knowledge-ascribing sentence in which ‘knows’ is either indexed to or relativized to an epistemic standard. So each such proposition is expressible and evaluable in any context.

II. Contextualism’s strategy for explaining the lure of skeptical arguments and resolving skeptical paradoxes requires that epistemologists not be cognizant of the previous point. If the relevant propositions are spelled out, then there is nothing to be confused about. And contextualism also imputes confusion to ordinary knowledge attributors: they implicitly take the standards for knowledge to vary with context and yet unwittingly take it be one and the same thing, knowledge, that is at issue from context to context. So contextualism is a strong error theory.

III. The contextualist strategy does not do justice to skeptical arguments. However cogent or fallacious such arguments may be, they purport to show that *ordinary* knowledge attributions are generally false. Skeptics argue not merely that we don’t have empirical knowledge by the highest standards but that we don’t have it at all (or at least not very much of it), even by ordinary standards. How assiduously people *apply* epistemic standards may vary from context to context, but the skeptic denies that the standards themselves come in various strengths. However loosely people ordinarily apply them, they are always highly demanding.

IV. In relying on skeptical arguments, skeptical invariantism implausibly attempts to draw semantic conclusions from epistemological considerations. Whatever the merits of these arguments, however, skeptical invariantism need not, and Unger’s version of it does not, rely on warranted assertibility maneuvers, lame or sophisticated. It stands or falls on the strength of skeptical arguments themselves. In my view it falls.

V. Moderate invariantism also does not have to rely on warranted assertibility maneuvers. The distinction between truth conditions and warranted assertibility conditions is a red herring, for the examples that contextualists use to motivate their thesis do not really provide evidence that ‘know’ is context-sensitive and that the truth conditions of knowledge-ascribing sentences can vary with the context in which they are uttered. What varies, rather, is the attributor’s threshold of confidence. In the problem cases, either a practical consideration or an overly demanding epistemic reason raises that threshold and leads the attributor to demand more evidence than knowledge requires.

VI. Contextualists agree, despite differences in formulations, that the varying standards governing [knowledge] attributions reflect what sorts of counterpossibilities need to be considered and eliminated. I suggest an entirely different way in which the consideration or non-consideration of counterpossibilities is relevant to having knowledge. In forming beliefs and seeking knowledge, we rely on our reliability to think of and thereby consider such possibilities when and only when they are worth considering. To the extent that we can trust our ability to know when there are no further counterpossibilities epistemically worth considering, we don't have to consider additional ones in order to be justified in treating them as not worth considering. This applies equally when we attribute knowledge to someone else.

Nobody disputes that how strictly we apply 'know', like a whole host of other words, varies with the context. But the way in which this is indisputable does not help the contextualist. As everyone knows, people use words with varying degrees of strictness and looseness, as I just did with 'everyone'. However, the fact that people do this does not show that the words themselves have semantic contents that come in various degrees. That may be true in the case of vague terms, but most contextualists do not claim that 'know' or 'knowledge' is vague. For all that the data about knowledge attributions show, it could well be that we often attribute knowledge to people who don't have it and often resist attributing it to people who do have it. Sometimes we speak casually, e.g., because we're interested in the answer to a certain question, hence in who has the answer, rather than in whether their true belief about the answer qualifies as knowledge. And sometimes we're extra cautious, say because of the stakes, and thus don't make up our own minds about the answer until we have obtained a second opinion, checked out possible sources of error, or otherwise confirmed the answer to our satisfaction. We can't attribute knowledge to someone, even if they have it, when we ourselves have doubts or worries about the truth of the proposition in question. Either for practical reasons or on dubious skeptical grounds, we sometimes demand more of knowledge than it requires.³⁹

³⁹ Many thanks to Jessica Brown, Ray Elugardo, John MacFarlane, Patrick Rysiew, and Jonathan Schaffer for very helpful comments and suggestions, which helped me forestall certain misunderstandings and meet certain objections.

References

- Austin, J. L. (1961). 'Other Minds', ch. 3 of *Philosophical Papers* (Oxford: Oxford University Press), 44-84.
- Bach, Kent (1984). 'Default Reasoning: Jumping to Conclusions and Knowing When to Think Twice', *Pacific Philosophical Quarterly* 65: 37-58.
- _____ (1985). 'A Rationale for Reliabilism', *The Monist* 62: 248-63. Reprinted in F. Dretske and Sven Bernecker (eds.), *Knowledge* (Oxford: Oxford University Press 2000), 199-213.
- _____ (1999). 'The Semantics-Pragmatics Distinction: What It Is and Why It Matters', in Ken Turner (ed.), *The Semantics-Pragmatics Interface from Different Points of View*. (Oxford: Elsevier), 65-84.
- _____ (2000). 'Quantification, Qualification, and Context', *Mind & Language* 15: 262-283
- _____ (2001). 'Speaking Loosely: Sentence Nonliterality', in P. French and H. Wettstein (eds.), *Midwest Studies in Philosophy*, vol. 25: *Figurative Language* (Oxford: Blackwell), 249-263.
- _____ (2002). 'Seemingly Semantic Intuitions', in J. Keim Campbell, M. O'Rourke, and D. Shier (eds.), *Meaning and Truth* (New York: Seven Bridges Press), 21-33.
- _____ (forthcoming). 'Context *ex Machina*', in Z. Szabó (ed.), *Semantics vs. Pragmatics* (Oxford: Oxford University Press).
- Brown, Jessica (forthcoming). 'Contextualism and Warranted Assertibility Manoeuvres'.
- Cappelen, Herman and Ernest Lepore (2003). 'Context-Shifting Arguments'. *Philosophical Perspectives* 17.
- Cohen, Stewart (1988). 'How to be a Fallibilist', *Philosophical Perspectives* 2: 91-123.
- _____ (1999). 'Contextualism, Skepticism, and the Structure of Reasons', *Philosophical Perspectives* 13: 57-89.
- _____ (2001). 'Contextualism Defended: Comments on Richard Feldman's 'Skeptical Problems, Contextualist Solutions'', *Philosophical Studies* 103: 87-98.
- DeRose, Keith (1992). 'Contextualism and Knowledge Attributions', *Philosophy and Phenomenological Research* 52: 913-929.

- _____ (1995). 'Solving the Skeptical Problem', *Philosophical Review* 104: 1-52.
- _____ (1999). 'Contextualism: An Explanation and Defense', in J. Greco and E. Sosa (eds.), *The Blackwell Guide to Epistemology* (Oxford: Blackwell Publishers), 187-205.
- _____ (2000). 'Now You Know It, Now You Don't', *Proceedings of the Twentieth World Congress of Philosophy*, Vol. V, Epistemology (Bowling Green, Ohio: Philosophy Documentation Center): 91-106.
- _____ (2002). 'Assertion, Knowledge and Context', *Philosophical Review* 111: 167-203.
- _____ (forthcoming). 'Single Scoreboard Semantics', *Philosophical Studies*.
- Dretske, Fred (1981). 'The Pragmatic Dimension of Knowledge', *Philosophical Studies* 40: 363-78.
- Feldman, Richard (1999). 'Contextualism and Skepticism', *Philosophical Perspectives* 13: 91-114.
- _____ (2001). 'Skeptical Problems, Contextualist Solutions', *Philosophical Studies* 103: 61-85.
- Goldman, Alvin (1976). 'Discrimination and Perceptual Knowledge', *Journal of Philosophy* 73: 771-91.
- Harman, Gilbert and Brett Sherman (forthcoming). 'Knowledge, Assumptions and Lotteries', *Philosophical Issues*.
- Hawthorne, John (forthcoming). *Knowledge and Lotteries* (Oxford: Oxford University Press).
- Heller, Mark (1999a). 'The Proper Role of Contextualism in Anti-Luck Epistemology', *Philosophical Perspectives* 13: 115-29.
- _____ (1999b). 'Relevant Alternatives and Closure', *Australasian Journal of Philosophy* 77: 196-208.
- Hofweber, Thomas (1999). 'Contextualism and the Meaning-Intention Problem', in K. Korta, E. Sosa, J. Arrazola (eds.), *Cognition, Agency and Rationality* (Dordrecht: Kluwer), 93-104.
- Kaplan, David (1989). 'Demonstratives', in J. Almog, J. Perry, and H. Wettstein (eds.), *Themes from Kaplan* (Oxford: Oxford University Press), 481-563.
- Kennedy, Christopher (forthcoming). 'Towards a Grammar of Vagueness'.

- Klein, Peter (2000). 'Contextualism and the Real Nature of Academic Skepticism', *Philosophical Issues* 10: 108-16.
- Kompa, Nikola (2000). 'The Context Sensitivity of Knowledge Ascriptions', *Grazer Philosophische Studien* 64: 1-18.
- Kornblith, Hilary (2000). 'The Contextualist Evasion of Epistemology', *Philosophical Issues* 10: 24-32.
- Lasersohn, Peter (1999). 'Pragmatic Halos', *Language* 75: 522-551.
- Lewis, David (1996). 'Elusive Knowledge', *Australasian Journal of Philosophy* 74: 549-67.
- Ludlow, Peter (this volume), 'Linguistic Tests for Contextualism in Epistemology'.
- Neta, Ram (2003). 'Skepticism, Contextualism, and Semantic Self-Knowledge', *Philosophy and Phenomenological Research* 67: 396-411.
- Nicolle, Steve and Billy Clark (1999). 'Experimental Pragmatics and What Is Said: A Response to Gibbs and Moise', *Cognition* 69: 337-354.
- Nichols, Shaun, Stephen Stich & Jonathan Weinberg (2003), 'Meta-Skepticism: Meditations on Ethno-Epistemology', in S. Luper (ed.), *The Sceptics* (Aldershot, U.K.: Ashgate Publishing), 227-247.
- Price, H. H. (1932). *Perception* (London: Methuen).
- Rysiew, Patrick (2001). 'The Context-Sensitivity of Knowledge Attributions', *Noûs* 35: 477-514.
- _____ (forthcoming). 'Speaking of Knowing'.
- Schaffer, Jonathan (forthcoming). 'From Contextualism to Contrastivism in Epistemology', *Philosophical Studies*.
- Schiffer, Stephen (1996). 'Contextualist Solutions to Scepticism', *Proceedings of the Aristotelian Society* 96: 317-33.
- Sosa, Ernest (2000). 'Skepticism and Contextualism', *Philosophical Issues* 10: 1-18.
- Stanley, Jason (forthcoming). 'On the Linguistic Basis for Contextualism', *Philosophical Studies*.
- Unger, Peter (1971). 'A Defense of Skepticism', *Philosophical Review* 80: 198-218.

_____ (1975). *Ignorance: A Case for Scepticism* (Oxford: Clarendon Press).

_____ (1984). *Philosophical Relativity* (Minneapolis: University of Minnesota Press).

_____ (1986). 'The Cone Model of Knowledge', *Philosophical Topics* 14: 125-78.

Williamson, Timothy (2000). 'Scepticism and Evidence', *Philosophy and Phenomenological Research* 60: 613-28.