

# Knowledge in and out of context

KENT BACH

kbach@sfsu.edu

## *Contents*

1. Making it explicit: the trouble with contextualism's anti-skeptical strategy
2. Facing up to skepticism
3. An invariantist take on the Bank and Airport cases
4. Practical interests and knowledge
5. The significance of salience
6. Conditional knowledge and conditional attributions
7. Summing up

Note: This paper is a revised version of "The emperor's new 'knows'." Sections 4 and 6 are completely new – §4 questions the Subject-Sensitive Invariantist take on contextualist examples, and §6 suggests that the reason many ordinary knowledge attributions seem true is that they either are implicitly attributions of conditional knowledge or are themselves implicitly conditional. The other sections are revised to various degrees, but if you've read "The emperor's new 'knows'," §4 and §6 should make sense on their own.

## Knowledge in and out of context

KENT BACH

We can be willing in one context to attribute a bit of knowledge that we wouldn't attribute and might even deny in another, especially a context in which we're stumped by a skeptical argument. Apparently, our standards for knowledge sometimes go up, sometimes way up. How can this be? By claiming that the very contents of knowledge-ascribing sentences vary with contexts of use, epistemic contextualism offers one explanation. I will offer another. According to contextualism, variation in standards is built into this claimed variation in contents. According to me, the contents of knowledge attributions are invariant. The variation is in what knowledge attributions we're willing to make or accept. Sometimes our standards are too strong, sometimes they're too weak, and sometimes they're just right.

Contextualism aims both to acknowledge and to escape the force of skeptical arguments. Consider these seemingly paradoxical reflections of David Lewis's:

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)

Contextualism is an ingenious way out: what counts as knowledge varies with the context in which it is attributed. More precisely, what so varies is the relation is expressed by the word 'knowledge' (or 'know'). So, depending on the context in which a sentence of the form 'S knows that p' is used, different relations between S and p are attributed.<sup>1</sup>

At first taste this idea is not easy to swallow. Despite the age-old difficulty of defining propositional knowledge, it still seems that the truth of a statement made using a sentence of the form 'S knows (at t) that p' depends at least on whether it is true that p, whether S believes that p, and whether S's belief that p is adequately supported *and* that

it does not at all depend on any facts about the person making the statement or on the context in which he is making it.<sup>2</sup> Why should it? Intuitively, it does not seem that the verb ‘knows’ expresses different relations in different contexts. So, for example, it seems that there is a unique proposition expressed by the sentence ‘Eminem knew on December 31, 1999 that Detroit was in Michigan’, namely, the proposition that Eminem knew on December 31, 1999 that Detroit was in Michigan. There do not seem to be even two different relations expressed by ‘know’, such that on that date Eminem knew<sub>1</sub> this but didn’t know<sub>2</sub> it. In this respect ‘know’ is very different from other allegedly context-sensitive terms, such as ‘poor’, ‘tall’, and ‘flat’.<sup>3</sup> With those terms, it takes just a little reflection on usage to be convinced that they can each express different properties. You don’t need to be presented with arguments to appreciate this, but with ‘know’ you do (assuming you are ultimately persuaded by contextualist arguments). So, at the very least, ‘know’ is not a typical context-sensitive term.<sup>4</sup>

Leaving that detail aside, notice that contextualism about ‘S knows that p’ does not concern the context of S (assuming the attributor is distinct from S). S’s circumstances do not affect the content of a knowledge attribution, but they can affect the truth-value (for example, S may be in a Gettier situation). Indeed, a new rival to contextualism, known as Subject-Sensitive Invariantism, goes so far as to suggest that practical importance to the subject, and perhaps even mere salience to him, can affect the truth-value of a knowledge attribution. Another recent view, put forward by John MacFarlane (2005), is that the truth-value of a knowledge attribution is not absolute but relative. Here there is no suggestion that the content of the knowledge attribution varies but only that its truth-value does, and that it varies not with the context of utterance but rather with the context of assessment. These two views raise interesting questions in their own right, but it is only contextualism that implies that the *content* of a knowledge attribution can vary, specifically with facts about the context in which it is made.<sup>5</sup>

In what follows, I will (1) explain why the contextualist anti-skeptical strategy is based on an illusion, that the propositions expressed by knowledge-ascribing sentences are themselves context-bound, (2) point out also that contextualism underplays the significance of skepticism, (3) offer a straightforward alternative to the contextualists’

interpretation of what's going on in their pet examples, (4) argue that the special features of their examples, practical interests and salience, do not directly affect people's epistemic position, (5) point out that, nevertheless, the salience of a possibility to a person can provide evidence for its relevance and that its nonsalience can function as tacit evidence for its irrelevance, and (6) suggest that what contextualists assume to be categorical knowledge attributions may really be attributions of conditional knowledge or else conditional attributions of knowledge. Although much of the following discussion is rather critical, I think the issues raised by contextualism are highly instructive and the challenges it poses are well worth confronting. So the discussion will also be constructive.

### **1. Making it explicit: the trouble with contextualism's anti-skeptical strategy**

Suppose that 'knows' is context-sensitive, as contextualists maintain. Even so, it would not follow that the propositions expressed, relative to contexts, by sentences of the form 'S knows that p' are themselves context-bound. This crucial point is generally overlooked or at least not highlighted by contextualists. The point bears repeating: even if contextualism were true, so that a given knowledge-ascribing *sentence* could express various propositions in various contexts, those *propositions* would not themselves be context-bound. To the contrary, each such proposition could be expressed by a more elaborate knowledge-ascribing sentence in which 'knows' is explicitly indexed or relativized to whatever it is that is supposed to vary with the context. Suppose it is an epistemic standard E.<sup>6</sup> Then *each* such proposition would be expressible in *any* context by the same more elaborate, context-insensitive sentence of the form 'S knows-by-E (at t) that p'.<sup>7</sup> It just has to be spelled out.

Contextualists neglect this crucial point when they implement their strategy for explaining the lure of skeptical arguments and resolving skeptical paradoxes. This neglect fosters the false impression that somehow the different propositions (allegedly) expressible (at a given time) by a given sentence of the form 'S knows that p' cannot be considered in the same context. Contextualists exploit this impression when they compare knowledge attributions made in ordinary contexts with those made in skeptical or

otherwise more stringent contexts.<sup>8</sup> As I'll now explain, once it is clear that distinct propositions, say the propositions that Eminem knew-by-E<sub>1</sub> that Detroit was in Michigan and the proposition that Eminem knew-by-E<sub>2</sub> Detroit was in Michigan, *can* be considered in the same context, the contextualist strategy for both explaining and neutralizing the lure of skeptical arguments loses its bite.

What is the contextualist strategy?<sup>9</sup> Consider the claim that Moore knew he had at least one hand, which to most people seems undeniable. The skeptical paradox arises from the fact that people are also moved by skeptical arguments, such as those based on far-fetched but seemingly hard-to-rule-out possibilities. One such possibility is that Moore was a bodiless (hence handless) brain in a vat, in short a BBIV. In light of this possibility the skeptic proposes an argument like this:

*Skeptical Argument*

If Moore knew he had at least one hand, then Moore knew that he wasn't a BBIV.

Moore didn't know that he wasn't a BBIV.

---

Moore didn't know that he had at least one hand.

Since this argument is formally valid, the only direct way to rebut it is to reject one of its premises, but that is not the contextualists' approach. Their strategy is subtler, aiming to expose an alleged equivocation on 'know'. The equivocation is not across steps of one argument but across contexts. That is, the above wording of what appears to be a single Skeptical Argument masks different arguments, each expressible by the same sentences used in different contexts. Each such argument is valid, but at most one is sound, and that one has no drastic skeptical consequences. The above form of argument is sound only if 'knows' is applied according to the highest epistemic standards, but those standards are operative only in a "skeptical" context, a context in which skeptical considerations and far-fetched possibilities (such as the BBIV) are raised. In an ordinary, less demanding context, where 'knows' expresses a weaker relation, a different argument is on the table.<sup>10</sup> The argument in that context is unsound, because its second premise is false.<sup>11</sup>

So contextualists respect skeptical intuitions as well commonsense, Moorean ones, but deny that these intuitions conflict. They thereby reject a more straightforward

assessment: a intuitions that give rise to skeptical paradoxes really do conflict and thus do not settle questions about as to the truth-values of knowledge attributions; rather, they comprise genuinely competing responses to skeptical considerations. Besides, people don't feebly capitulate to skeptical arguments like the one stated above but tend to vacillate. They go back and forth between finding the argument hard to refute and finding its conclusion hard to swallow. Their resistance to its conclusion cannot be written off as a Humean lapse back into an ordinary, non-skeptical context, for it occurs while skeptical possibilities are firmly in mind. So it seems that contextualists are not really entitled to rely on their intuition that knowledge attributions made in skeptical contexts are false.

But suppose they are entitled. In claiming that the conclusion of the Skeptical Argument in a skeptical context is compatible with the claim made in an ordinary context that Moore knew he had hands, the contextualist has to concede that people not privy to contextualism do not realize that the relation expressed by 'know' can shift with context.<sup>12</sup> After all, they are taken in by the skeptical argument, at least enough to take it seriously and without any sense that the subject has been changed. By allowing the skeptical argument to call into question the ordinary claim that Moore knew he had hands, they both mistakenly and unwittingly take it be one and the same thing, knowledge, which is at issue from context to context. So contextualism is a kind of error theory. To be sure, people will make and accept ordinary knowledge claims once they're removed from a skeptical context (does this require a bit of amnesia?), but that only compounds the error, since they do not realize that what they're now willing to accept as true is not the same proposition as the one they rejected in the skeptical context. So the contextualist claim that people's contrasting intuitions seem to but don't really conflict requires that people don't notice these shifts in contextual shifts in the relation expressed by 'knows'.

This is clear when we apply our earlier observation that even if standard knowledge-ascribing sentences are context-sensitive, the different propositions they supposedly express in different contexts are not themselves context-bound. We can make explicit two different arguments, one presented in an ordinary context and one in a skeptical context, which both take the form of the Skeptical Argument, so that both may be stated in any

context, including this one. Let knowledge-by-ordinary-standards be expressed by ‘knows<sub>o</sub>’ and knowledge-by-skeptical-standards by ‘knows<sub>s</sub>’. Then the two arguments take the following forms:

*Skeptical Argument in an ordinary context*

If Moore knew<sub>o</sub> he had at least one hand, then Moore knew<sub>o</sub> that he wasn’t a BBIV.

Moore didn’t know<sub>o</sub> that he wasn’t a BBIV.

---

Moore didn’t know<sub>o</sub> that he had at least one hand.

and

*Skeptical Argument in a skeptical context*

If Moore knew<sub>s</sub> he had at least one hand, then Moore knew<sub>s</sub> that he wasn’t a BBIV.

Moore didn’t know<sub>s</sub> that he wasn’t a BBIV.

---

Moore didn’t know<sub>s</sub> that he had at least one hand.

There is now no question of unwittingly shifting from one context to another and thereby confusing the Skeptical Argument in a skeptical context with the Skeptical Argument in an ordinary context. We can just look at the two arguments together, regardless of our context, and evaluate the two of them. Now assume that contextualists’ intuitions about the truth-values are correct. Then the first argument is unsound, since its second premise is false. Moore *did* know<sub>o</sub> that he had at least one hand. But the second argument is sound, since its second premise is true; or at least the second argument will seem sound, and seem irrefutable, to anyone moved by skeptical considerations and unable to rebut its second premise. But the conclusion of the second argument is relatively innocuous, since it denies only that Moore *knew<sub>s</sub>* (knew by the most demanding standards) that he had at least one hand.

So the contextualist strategy for resolving the skeptical paradox loses its bite once we use different phrases, ‘knows by ordinary standards’ and ‘knows by skeptical standards’ or, as above, abbreviate these phrases with the help of different subscripts. Since now everything is now out on the table, we can intuit without conflict that there are lots of things we can know by ordinary standards even if we can’t know them by skeptical

standards — assuming, of course, that there are different knowing relations. Later (in section 3) we will see that there is no reason to concede this, but first let's look further at contextualism's attempt to marginalize skepticism.

## 2. Facing up to skepticism

In attempting to confine the plausibility of skeptical arguments to contexts in which far-fetched skeptical possibilities are raised, contextualism doesn't really do justice to those arguments, however cogent or fallacious they may be. Hilary Kornblith's (2000) title, "The Contextualist Evasion of Epistemology," is apt, for skeptical arguments purport to show that *ordinary* knowledge attributions are generally false.<sup>13</sup> The contextualist's attempt to marginalize these arguments by restricting them to skeptical contexts ignores the fact that skepticism denies that we have knowledge 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. 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, that is, to have the one property actually and ordinarily expressed by the word 'knowledge'. He is not proposing to reform its meaning. He is recommending that we keep using it to mean what it ordinarily means but use it much more carefully, even if it turns out rarely to apply.

In light of this, must the invariantist capitulate to skepticism? No, because invariantism is a semantic thesis about 'know', not a substantive theory of knowledge. Even so, the invariantist needn't be completely neutral and has a way of countering the lure of skeptical arguments. When we are 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, being 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 what sort of world we're being



plunked into (or, as the skeptic would rather put it, about what sort of world we're actually in). Since each of these possible worlds is consistent (let's assume) 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.<sup>14</sup> 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 which one we're in or whether it's at all like the world we think we're in.

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 undetectable and uncorrectable error, at least not if knowing this requires going out and verifying that we're not. True, a skeptical scenario would seem no less absurd if we 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, such that a great many of our expectations and intentions appear to be fulfilled, is to be in a world in which there is plenty of knowledge to be had. Although our evidence does not deductively eliminate the possibility that we live in some sort of fanciful skeptical world, skeptical arguments offer no reason to suppose that we do. And, if only because of the intractable computational complexity required to maintain such a world, these worlds are far more improbable than Hume thought miracles to be.

Whether or not this is even the beginning of a good answer to the skeptic, at least it doesn't accuse of skeptic of changing the subject. It does concede to the skeptic that if we were BBIVs, we would take our epistemic position with respect to things we claim to

know to be as good as we actually think it is, and no doubt the skeptic would pounce on that. However, the skeptic has a similar basis for responding to the contextualist. He could point out that if contextualists were BBIVs, they would still believe that ‘know’ is context-sensitive and base this belief on the same intuitions that they actually rely on. But in that case the intuition that many of our ordinary knowledge attributions are true would be mistaken. Yet ‘know’ would mean the same thing as it does in what we take to be the actual world.

At any rate, there is one thing that everyone can agree on: the sorts of possibilities used to support skeptical arguments have to be so sweeping that no kind or amount of further inquiry can eliminate them. It’s inherent in these possibilities, as standardly laid out, that any inquiry aimed at ruling them out is vulnerable to the same or similar possibilities, so that no progress in ruling them out can ever be made. In this important respect, they are unlike the sorts of error possibilities we take seriously in everyday life, which we can and do take effective measures to rule out. However, contextualists do not rely solely on the contrast between ordinary and skeptical contexts, between contexts in which we take much for granted and those in which everything is up for grabs. It’s a key part of their strategy to invoke a different contrast, between ordinary undemanding contexts and merely *relatively* demanding ones, where people feel they need to eliminate realistic possibilities of error.

### **3. An invariantist take on the Bank and Airport cases**

Whether skepticism is refutable or merely indefensible, you don’t have to be a skeptic to be an invariantist, certainly not for the purpose of rebutting contextualism. From here on I will go along with contextualists’ assumption that a good many of the knowledge attributions we make are true, at least in ordinary contexts. However, I will offer a moderate (nonskeptical) invariantist alternative to their interpretation of the high-standards versions of their pet examples.

Their most well known examples are Keith DeRose’s Bank case (1992: 913) and Stewart Cohen’s Airport case (1999: 58). I will assume that the reader is familiar with the details. In each case we are asked to make an intuitive comparison between a knowledge

attribution made in a normal, low-standards context and one made in a high-standards but non-skeptical context. It is assumed that the subject has the same evidence in both the low- and the high-standards versions of each case and that the attributor's evidence is essentially the same (identical or at least of equal strength) as the subject's. Even so, according to contextualists' intuitions, whereas the knowledge attribution made in the normal, low-standards context is true, the one made (using the same knowledge-ascribing sentence) in the high-standards context is false.<sup>15</sup>

Contextualists contend that these contrasting intuitions do not really conflict and can both be correct, because 'know' is context-sensitive so that the truth-conditions of knowledge-ascribing sentences can vary with the context in which they are used. I think there's a simpler, more straightforward interpretation, according to which the two intuitions really do conflict, just as they seem to, hence cannot both be accepted as correct.<sup>16</sup> On this interpretation, what varies with context is the attributor's threshold of confidence. In the high-standards context, it is raised: either a practical concern or (in skeptical cases) an excessive epistemic demand leads the attributor to demand more evidence than knowledge requires. In contexts where special concerns arise, whether practical or skeptical, what varies is not the truth-conditions of knowledge attributions but the knowledge attributions people are prepared to make. What goes up is not the standard for the truth of a knowledge attribution but the threshold of confidence the attributor has to reach to make the attribution.

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 the Chicago stopover 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 the contextualist, that's because it isn't true that Smith knows this (by high standards). 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. How could this be?

Mary is making a mistake, albeit a very understandable one. It is because of her own doxastic situation that she does not say, “Smith knows that the plane will stop in Chicago,” and goes so far as to assert its negation. Because she is not sure his itinerary is reliable, she herself does not believe, at least not confidently, that the plane will stop in Chicago. So 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 it still needs to be verified. Not only can’t Mary very well assert that Smith knows that the plane will stop in Chicago, 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.

Now what is essential here is not the attributor’s lack of settled belief but her raised threshold for (confidently) believing.<sup>17</sup> Before believing the proposition in question, at least with the confidence and freedom from doubt necessary for knowing, the attributor demands more evidence than knowledge requires. Look at what happens, in the high-standards version of the Airport case, when Mary checks further until she is satisfied 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.

One’s threshold for (confidently) believing a proposition is a matter of what one implicitly takes to be adequate evidence. 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 in a position to know it is enough to keep one from believing it (confidently) 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 knowledge attributions. However, my account does not imply a shift in their contents. I have simply pointed out a constraint on 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. Considering that attributing knowledge that *p* requires confidently believing that *p*, I am suggesting that willingness to attribute knowledge tracks not the standards on the truth of a knowledge attribution but one's threshold of doxastic confidence. In the so-called high-standards cases, the attributor's confidence threshold goes up to the point that without additional evidence she implicitly, but mistakenly, thinks she is not in a position to know. In high-standards Airport and Bank cases, a special practical interest makes it more difficult than normal for the attributor to have a settled belief in the relevant proposition and, accordingly, raises her bar for attributing knowledge to someone else. Even so, that person 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. Does merely raising such a possibility, without making it plausible, turn a true knowledge attribution into a false one? Making it salient does not do that.<sup>18</sup> As Patrick Rysiew (2001) has shown, making it salient can affect only the assertibility of the knowledge attribution, because attributing knowledge pragmatically conveys that a newly raised possibility has been ruled out.<sup>19</sup> Raising plausible possibilities, on the other hand, indicates genuine doubts or worries on the attributor's part and may lower the audience's confidence level accordingly. Moreover, if these plausible possibilities are objective ones, in which case they 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, quite independently of the attributor's context. So the truth-condition and the truth-value of a knowledge attribution are not 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, as the result of their confidence threshold having been raised.

We are now in a good epistemic position to meet a challenge posed by DeRose for moderate invariantism in its effort to account for intuitions about ordinary and high-standards knowledge attributions. He writes,

[Invariantists] 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)

The problem here is that DeRose accepts the appearances at face value.<sup>20</sup> 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'.

Moderate invariantists should also reject the intuition that knowledge denials involving skeptical possibilities are true. The moderate invariantist should not concede that there is something right about the intuition that Moore does not know he is not a BBIV and that an utterance of 'Moore does not know that he is not a BBIV' is true, at least in an epistemic context. Rather, he should insist that Moore 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 Moore's epistemic situation would be no different. To be sure, Moore doesn't 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 Moore, 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.<sup>21</sup> 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 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 obviously an error theory, and so is contextualism. I gladly concede that as I have defended it moderate invariantism is an error theory too, but only in a minimal way. According to skeptical invariantism, the knowledge attributions that people ordinarily make are almost all false. According to contextualism, people commonly fail to recognize shifts in the contents of knowledge attributions and thereby sense contradictions that are not there; indeed, people are even confused about what they themselves mean when they use 'know' in different contexts. But the only sort of error that moderate invariantism (as I have defended it) 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 epistemic caution when it comes to believing things with confidence. But such caution has a practical rationale and is therefore not irrational. Sometimes it is reasonable to go beyond the call of epistemic duty.

#### **4. Practical interests and knowledge**

I have suggested that in the high-standards versions of the Bank and Airport cases, the attributor, having practical reasons for not accepting the proposition in question merely on the basis of the subject's evidence, understandably but mistakenly denies that the subject knows, even though (it is assumed on both sides) the subject's evidence is

adequate for knowledge. Given the attributor's practical interests, he deems the subject's evidence to be inadequate. But what about the *subject's* practical interests? Can these, or some other not obviously epistemically relevant aspect of the *subject's* context, affect the truth-value of a knowledge attribution?

That's the idea behind the new "Subject-Sensitive Invariantism." SSI denies that the content, the truth-condition, of a knowledge attribution is affected by the context of the attributor, but it proposes that the subject's context can matter.<sup>22</sup> SSI is not just another form of relevant alternatives theory. It concerns practical interests, not neutral epistemic reasons. SSI says that the subject's practical interests can make a difference as to whether or not the subject knows. That is, two people can be in identical evidential positions but different epistemic positions, such that one knows a certain thing while the other does not know the counterpart of that thing, and that's because of what is at stake.<sup>23</sup>

Underlying SSI is the crucial assumption that one can't know that p if one has reason to make sure that p. Of course, it is uncontroversial that one can't know that p if one has epistemic reason to make sure that p, but SSI goes further. It supposes one can't know that p just because one has practical reason to make sure.<sup>24</sup> But why suppose this? Or, insofar as we should suppose this, why should it require turning simple invariantism into the subject-sensitive variety? If it is true that p but it makes no difference to you whether or not p, then you are in a position to know that p if you have sufficient epistemic reason to believe it.<sup>25</sup> So how can having some practical interest in whether or not p keep you from being in a position to know that p? How can suddenly caring enough to be disposed to guard against possibilities that ordinarily wouldn't concern you deprive you of that knowledge?

How this is possible is obvious from the previous section's discussion of the Bank and Airport cases: your practical interest may lead you to want to make sure that p before you act on the supposition that it holds true. As a result, you don't yet believe that p, at least not with confidence, and wish to guard against certain possibilities of error. This means that you don't yet know that p. The reason for this is not that you have insufficient epistemic reason for believing that p but that you don't meet the doxastic condition on knowing. So cases of this sort do not support SSI.



A different situation arises if you do have a settled and confident belief that *p*. Even then, given the cost of being wrong you may think that you need to make sure that *p*, by ruling out certain possibilities of error that ordinarily would be too remote even to consider, much less bother with. But this sort of situation does not support SSI either. Rather, it shows that one can know that *p* even if one has reason to make sure that *p*, contrary to the crucial assumption underlying SSI. Again, this is because of the cost of being wrong. It is rational to check further for essentially the same reason that it is rational to make a small bet against something you know if the payoff is large enough.<sup>26</sup> Thinking you might be wrong is compatible with knowing. What is not compatible with knowing is having a specific reason for thinking you might be wrong. But there's nothing irrational about recognizing one's fallibility and the general possibility of error. It might be odd to say, "I know but I need to make sure," but in this case it makes sense. Similarly, it would be coherent for an observer of the situation to say, "He knows, but he ought to check further." That would ordinarily sound odd, but not in this case. So, contrary to the assumption underlying SSI, having practical reason to make sure that *p* is compatible with knowing that *p*. When it keeps one from knowing that *p*, this is only because it keeps one from meeting the doxastic condition on knowing.

What about salience? Can the mere fact that a counterpossibility to *p* is brought up, or comes to mind, keep one from knowing that *p*? Can this do for a subject's epistemic position what contextualists think it can do for an attributor's? Hawthorne briefly considers the possibility that salience "destroys" belief or at least weakens it enough to destroy knowledge, but he rightly sees that this is no problem for simple invariantism, since it "hardly takes us beyond the factors traditionally adverted to in accounts of knowledge, given the centrality of the belief condition to standard accounts" (2004: 173).<sup>27</sup> Nevertheless, there is an epistemic role for salience (of a possibility of error), but one that also comports with simple invariantism. It is a role that Hawthorne himself considers but does not pursue, that "salience destroys knowledge by providing counterevidence" (2004: 172). As we will see next, sometimes it does, but only sometimes.

## 5. The significance of salience

There is more to be said about the interaction between one's doxastic state and one's epistemic position and about how this bears on 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 (these can be alternatives to *p* or threats to the basis for one's belief that *p*). There is plenty of dispute about how best to formulate this, especially if one rejects the skeptic's insistence 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 some of David Lewis's reflections in this regard.

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 actually in it. So-called skeptical “hypotheses” are really 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). Any possibility compatible with the experience (with its having the content that it has) is not eliminated.<sup>28</sup> 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. But how can they be properly ignored without first being eliminated? For 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 a way around the Rule of Attention.

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. Suppose you're looking at a zebra. If it doesn't occur to you that there might be some cleverly painted mules in the vicinity, you can't very well rule out that possibility. Does this keep you from being in a 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? Considering how far-fetched it is, can't you just dismiss it? The mere fact that the thought of it occurs to you shouldn't make it harder to dismiss. Or should it?

Offhand, it might seem that whether or not the thought of a certain possibility occurs to us has no epistemic significance. 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 is a distraction and, if chronic, would be a nuisance, but that would be that. We resist making what John Hawthorne calls "anxiety-provoking" inferences (2004: 161) because we take our anxiety, hence its source, seriously (of course, some people are more cautious than others, and each of us is more cautious in some areas than in others). 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 is *not* worth considering.<sup>29</sup>

If our cognitive processes are operating well, generally the thought of a possibility contrary to something we're inclined to believe occurs to us only if it's a realistic

possibility, not a far-fetched one (this is, of course, a matter of degree). This happens, for example, when we recall a bit of information but then think of a plausible alternative. I recall that it was Senator Hruska of Nebraska who said, regarding a Nixon judicial nominee who was widely regarded as mediocre, that mediocre people need to be represented on the Supreme Court. But then I wonder if it wasn't Senator Dirksen of Illinois who said that. Then I think, if it was Dirksen, who was much more well known, I wouldn't have thought it was someone as obscure as Hruska. So I rule out Dirksen and settle on Hruska. We can always conjure up wild possibilities, as in flights of skeptical or paranoid fantasy, but when we're trying to remember some bit of information, trying to identify what we're perceiving, or are engaged in normal inquiry, we take into account only those sorts of possibilities of error that sometimes arise in situations of the sort we're concerned with, and we do so because we rely on our ability to think of them when and only when they are worth taking seriously. 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"! Should it occur to us we may of course find reason to dismiss it.

Part of cognitive competence in a given area consists in reliably thinking of possibilities of error when they are worthy of consideration and at knowing when to look further before settling into a belief — and, in being able to rely on this reliability. Not only that, we implicitly take the nonoccurrence of the thought of a possibility of error to be evidence, albeit highly defeasible, that this possibility is *not* worth considering. After all, we cannot *explicitly* take the *nonoccurrence* of the thought as the evidence that it is, for that would entail thinking of the very possibility in question. We cannot explicitly weigh the evidence that the nonoccurrence of the thought provides, at least not at the time (sometimes, in retrospect, we reason that if something was a realistic possibility, it would have occurred to us). Instead, we rely on the reliability of our tendency to think of error possibilities when and only when they're worth considering. We implicitly assume that if a given possibility of error deserved consideration, we would have thought of it.

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 has practical reasons or bad skeptical reasons for ruling out, 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 section 3, this is enough to put him in the position of refraining from asserting that the subject knows and even of falsely denying that the subject knows. As we can now see, what does this, and what can unsettle the attributor's own belief, is the consideration of a possibility that epistemically is not worth considering.

## **6. Conditional knowledge and conditional attributions**

In section 3, I offered an invariantist interpretation of the Bank and Airport cases. It does not accept the contextualist intuition that when the standards go up, knowledge attributions that would ordinarily be true are not true and the corresponding knowledge denials are true. Instead, it proposes to explain why we are willing to make the knowledge attribution in one context but not in another. Now, partly relying on some ideas from the last two sections, I would like to suggest, without endorsing, two other possible ways of dealing with contextualist data.

Perhaps knowledge is not as easy to come by as we have been assuming. Moderate invariantism does not make a specific claim as to how easy, but in going along with contextualist stipulations I did assume that in the low-standards versions of the Bank and Airport cases the subject really did know. This might be stretching things a bit. You hardly have to be a skeptic to acknowledge that there are plenty of down-to-earth sources of error that we don't guard against most of the time. So perhaps much of what passes for categorical knowledge is really conditional knowledge. That is, when it seems that a person knows that  $p$ , it may really be that the person knows that  $p$ , if  $C$ , where  $C$  is a normality condition that is implicitly taken for granted.<sup>30</sup> So, for example, if you've left

your car in your driveway, you know that it is in your driveway, provided it hasn't been moved since. You implicitly assume that your car hasn't been stolen, towed away, or hit by a runaway truck, etc.<sup>31</sup> And, if your knowledge about your car's whereabouts is conditional, perhaps so is Smith's knowledge that his plane will stop in Chicago. Perhaps Smith doesn't categorically know this. Perhaps what he knows is that his plane will stop in Chicago provided the itinerary is accurate, the flight won't be diverted from Chicago because of bad weather (the most likely counterpossibility), and the plane won't crash or be hijacked (etc.). Being not a theory of knowledge but a view about the semantics of 'know' (and its cognates and synonyms), moderate invariantism can be neutral about how much of what people are said to know they really do know and about how much of this is really conditional in content.

Now it might seem extreme to suppose that most of what ordinarily passes for categorical knowledge is really conditional in content. So consider a more modest suggestion. Yes, this knowledge is categorical in content, but its being knowledge is contingent on a normality condition.<sup>32</sup> So whereas the previous idea was that much of what passes for categorical knowledge (that p) is really conditional in content (that p, if C), here the idea is that the knowledge that p is conditional.<sup>33</sup> If so, then many of our ordinary knowledge *attributions*, insofar as they can be deemed true, are best interpreted as conditional on the implicit assumption that things are normal. For example, an utterance of 'Smith knows that the plane will stop in Chicago' is not categorically true but true only if understood as a loose, casual way of saying, 'Provided things are normal, Smith knows that the plane will stop in Chicago'. Sometimes, the attributor may even intend this normality condition to be understood, even though it is not part of the explicit, literal content of the attribution.<sup>34</sup>

Each of these suggestions offers an alternative to the contextualist way of explaining how the intuitions that generate the skeptical paradox do not really conflict. Even if the contextualist is correct that they don't, this might not be a matter of 'know' expressing different relations in high-standards than in low-standards contexts. On the first suggestion, there is indeed a difference in content between knowledge attributions made using the same knowledge-ascribing sentence in the two contexts, but this difference

consists instead in what knowledge is attributed. In low-standards contexts it is conditional knowledge (that p, if C), even though the sentence ('S knows that p') does not make that condition explicit, whereas in high-standards contexts it is unconditional knowledge that p. Alternatively, perhaps the "low-standards" attribution is itself conditional and "high-standards" one is not. And allowances could be made for intermediate cases by different degrees of strengthening of the condition — either in what is known or in the condition on knowing it.

One complication arises from questions raised by the extended version of what John Hawthorne (2004) calls the "lottery" puzzle. The puzzle, which is not limited to lotteries, is to explain how one can know an "ordinary" proposition but fail to know a related "lottery" proposition, that is, the highly likely proposition that a certain counterpossibility doesn't obtain.<sup>35</sup> The ordinary proposition could be, for example, that your car is in your driveway and the lottery proposition that your car has not been stolen (since you parked it). The puzzle is both to explain and to reconcile the intuitions (preferably without giving up closure) that you know the first but don't know the second and, in particular, to account for the fact that even if do know that your car is in your driveway, you are not entitled to infer from this that your car has not been stolen. The complication, as illustrated by this example, arises from the conflict between two intuitions: that you know your car is in your driveway, and, that you do not know it has not been stolen.<sup>36</sup> Taking the intuitions at face value suggests that, somehow, it is easier to know that your car is in your driveway than that it has not been stolen. This suggestion seems very odd, especially assuming closure, since it seems that you could not know that your car is in your driveway unless you know that it hasn't been stolen.

With this complication in mind, let's apply the two ideas offered above to this complication. The first idea was that much of what passes for categorical knowledge is really conditional in content. In this example, what you know is not that your car is in your driveway but merely that it is in your driveway if it has not been stolen (pretend that there are no other plausible error possibilities). Insofar as it is reasonable to act on what you know, it is reasonable to act on the supposition that your car is in your driveway even though you don't know that categorically, provided it is reasonable to suppose that your

car has not been stolen. For example, if you soon have to get somewhere by car and you're not far from home, it would be reasonable to walk home with the intention of taking your car. However, it would be unreasonable, in order to reduce your premium, to cancel the theft coverage on your auto insurance policy on the grounds that you know that your car is in your driveway (and generally know the whereabouts of your car). This is partially explained by the fact that reasonably acting on the supposition that your car has not been stolen is not predicated on an outright conviction that your car has not been stolen. That proposition is not a lottery proposition with respect to itself. That is, even if knowledge that  $p$  if  $C$  can pass for knowledge that  $p$ , when it is reasonable to assume  $C$ , knowledge that  $C$  if  $C$  obviously cannot pass for knowledge that  $C$ .

Now let's apply the second of the above ideas. Saying that  $S$  knows that his car is in his driveway could be predicated on the condition that his car hasn't been stolen; but saying that  $S$  knows that his car hasn't been stolen couldn't be predicated on the condition that his car hasn't been stolen. Consider that when we evaluate a knowledge attribution and, indeed, when we consider making one and judge whether someone knows a certain proposition, we ask questions of this form: would  $S$  have believed that  $p$  even if not- $p$ ? As DeRose (1995) makes clear, the possibilities of error that we take into account can vary. In some contexts, and quite apart from the subject's situation, we consider more remote possibilities than in other contexts. That is, sometimes we implicitly assume much less than we do on other occasions, so that more is, so to speak, up for grabs. This is virtually forced upon us when we shift our focus from whether  $S$  knows that  $p$  to whether  $S$  knows that not- $e$ , where  $e$  is some possibility of error. For if we previously took for granted that not- $e$  when we were otherwise inclined to attribute to  $S$  knowledge that  $p$ , we can't very well take for granted that not- $e$  when  $e$  itself comes into consideration. So if that attribution were implicitly conditional on the supposition that not- $e$ , then naturally we would regard it as true so long as we implicitly assumed that not- $e$ . We would not so regard it, hence would not be inclined to make it, if we no longer implicitly assumed that not- $e$  and regarded this as unsettled. In making this shift, we haven't used 'know' to express a stricter relation than before. Rather, we have gone from presupposing more to presupposing less in regard to the satisfaction of the conditions for  $S$ 's knowing that  $p$ .



The moral of the last two paragraphs is that the data cited by contextualists as best explained by the supposition that knowledge attributions are context-sensitive might better explained in this way: a great many of the knowledge attributions that we are willing to make and to accept are either attributions of implicitly conditional knowledge or are themselves implicitly conditional (insofar as they can be taken to be true). Whether or not either of these ideas is right, they at least suggest alternative ways of accommodating contextualist data without resorting to a shifty semantics for ‘know’.

## 7. Summing up

Epistemic contextualism says that ‘knows’ expresses different relations and simple knowledge-ascribing sentences express different propositions in different contexts. However, as pointed out in section 1, even if this is correct the propositions thus expressed would not themselves be context-bound. Indeed, they could be made fully explicit by using more elaborate sentences that specify the operative standards (or whatever it is that is supposed to vary with context). Then it would be obvious that these propositions could be compared and evaluated in any context, thus eliminating the surreptitious semantic shiftiness on which contextualism bases its anti-skeptical strategy and, as suggested in section 2, neutralizing the contextualist attempt to marginalize skeptical arguments by confining them to skeptical contexts. Even if we cannot consider skeptical arguments in ordinary contexts, since entertaining far-fetched possibilities *ipso facto* puts us in a skeptical context, the propositions that make up these arguments and so the arguments themselves are not context-bound. Either they are sound or they are not, independently of the contexts in which we consider them. However cogent or fallacious these arguments may be, skeptics intend them to show that *ordinary* knowledge attributions are generally false.

Section 3 offered a straightforward invariantist interpretation of the data in the contextualist’s high-standards scenarios and in skeptical contexts. What goes up is not the truth-condition of the relevant knowledge attribution but the attributor’s threshold of confidence. In the problematic scenarios, either a practical consideration or an overly demanding epistemic reason raises that threshold and leads the attributor to demand more

evidence than knowledge requires. Accordingly, he can't coherently attribute knowledge to someone else even if they have it. Indeed, he must go so far as to deny that they know, since he regards their evidence as inadequate for them as it is for himself.

This observation is relevant also to Subject-Sensitive Invariantism, a recent rival to contextualism. A crucial assumption underlying SSI, as pointed out in section 4, is that one can't know that *p* if one has reason to make sure that *p*. But this is not because the stakes or salience have raised the evidential bar for knowledge. What is raised, rather, is the subject's confidence threshold for believing. Even if you're in a position to know that *p*, you don't, because you don't confidently believe that *p*. Before you can confidently do that you need to check further until you are satisfied that you have made sure that *p*.

Section 5 suggested a new twist on Lewis's observation that because what is salient or at stake can differ in different contexts in which a given knowledge-attributing sentence is used, there is concomitant variation in the sorts of error possibilities that need to be eliminated and in the ones that can be "properly ignored." I pointed out a common but commonly overlooked way in which the consideration or non-consideration of possibilities of error is relevant to having knowledge. In forming beliefs and seeking knowledge, we implicitly rely on our reliability to think of and thereby consider possibilities of error when and only when they are worth considering. Part of what makes beliefs justified is that the cognitive processes whereby they are formed and sustained, at least in areas of cognitive competence, are sensitive to realistic possibilities of error (so-called relevant alternatives). The very occurrence of the thought of such a possibility gives one *prima facie* reason to take it seriously. And the fact that a particular error possibility does not come to mind is evidence for its irrelevance. But that fact is evidence that we cannot explicitly consider, since to do so would be to bring the possibility to mind. To the extent that we can trust our ability to know when there are no further error possibilities epistemically worth considering, we don't have to consider them in order to be justified in treating them as not worth considering. This applies equally when we attribute knowledge to others.<sup>37</sup>

Section 6 pointed out that there are still other ways to deal with the apparently conflicting intuitions that motivate both contextualism and the straightforward alternative

suggested in section 3. It could be that the contextualist is correct in supposing that the relevant intuitions do not really conflict but wrong about why. Perhaps, when a given knowledge-ascribing sentence is used, there is a difference in content between knowledge attributions made in what contextualists describe as high-standards contexts and in low-standards contexts, but it is not a matter of ‘know’ expressing a different relation. Instead, perhaps what is attributed in low-standards contexts is conditional knowledge, even though the condition is not made explicit, whereas in high-standards contexts unconditional knowledge is attributed. Alternatively, perhaps “low-standards” attributions are themselves conditional and “high-standards” ones are not. Either way, the explanation for the difference in intuitive truth-value is that only in the high-standards cases does this condition have to be discharged. Knowledge may not be as difficult to achieve as skeptics imagine or as easy to come by as contextualists allow (in so-called low-standards contexts), but there is nothing shifty in the semantics of ‘know’.

## References

- 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.
- \_\_\_\_\_ (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.
- \_\_\_\_\_ (2005), "Context *ex Machina*," in Z. Szabó (ed.), *Semantics vs. Pragmatics* (Oxford: Oxford University Press), 15-44.
- Cohen, Stewart (1999), "Contextualism, Skepticism, and the Structure of Reasons," *Philosophical Perspectives* 13: 57-89.
- DeRose, Keith (1992), "Contextualism and Knowledge Attributions," *Philosophy and Phenomenological Research* 52: 913-929.
- \_\_\_\_\_ (1995), "Solving the Skeptical Problem," *Philosophical Review* 104: 1-52.
- \_\_\_\_\_ (2002), "Assertion, Knowledge and Context," *Philosophical Review* 111: 167-203.
- Fantl, Jeremy and Matthew McGrath (2002), "Evidence, Pragmatics, and Justification," *Philosophical Review* 111: 67 - 94.
- Feldman, Richard (1999), "Contextualism and Skepticism," *Philosophical Perspectives* 13: 91-114.
- \_\_\_\_\_ (2001), "Skeptical Problems, Contextualist Solutions," *Philosophical Studies* 103: 61-85.
- Harman, Gilbert and Brett Sherman (2004), "Knowledge, Assumptions and Lotteries," *Philosophical Issues* 14: 492-500.
- Hawthorne, John (2004), *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.
- Klein, Peter (2000), "Contextualism and the Real Nature of Academic Skepticism," *Philosophical Issues* 10: 108-16.
- Kornblith, Hilary (2000), "The Contextualist Evasion of Epistemology," *Philosophical Issues* 10: 24-32.
- Lewis, David (1996), "Elusive Knowledge," *Australasian Journal of Philosophy* 74: 549-67.
- John MacFarlane (2005), "The Assessment Sensitivity of Knowledge Attributions," in T. Gendler and J. Hawthorne (eds.), *Oxford Studies in Epistemology* (Oxford: Oxford University Press), 000-000.
- 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.
- Rysiew, Patrick (2001), "The Context-Sensitivity of Knowledge Attributions," *Noûs* 35: 477-514.
- Schaffer, Jonathan (forthcoming), "What Shifts?" "Thresholds, Standards, or Alternatives?"
- 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 (2004), "On the Linguistic Basis for Contextualism," *Philosophical Studies* 119: 119-146.
- \_\_\_\_\_ (forthcoming), "Context, Interest-Relativity, and Knowledge."
- Unger, Peter (1971), "A Defense of Skepticism," *Philosophical Review* 80: 198-218.
- Williamson, Timothy (2000), *Knowledge and its Limits*. Oxford: Oxford University Press.

## Notes

<sup>1</sup> For this reason, from a contextualist perspective it is misleading to speak of knowledge attributions without qualification. Although I will persist in using the phrase ‘knowledge attribution’, it would be more accurate, when discussing contextualism, to use more a cumbersome phrase, like ‘literal assertive utterance of a “knows”-ascribing sentence’.

<sup>2</sup> I will not dwell on the fact that contextualists are not very clear on the semantic nature of this alleged context sensitivity. Is it indexicality, vagueness, relativity, ambiguity, semantic underdetermination, or what? They sometimes casually speak of ‘knows’ as an indexical term, not that ‘know’ behaves like ‘I’ or ‘now’, but sometimes they liken ‘know’ to relational terms (‘local’, ‘enemy’) or to gradable adjectives (‘tall’, ‘flat’). Jason Stanley (2004) has forcefully argued that ‘know’ does not fit any of these models. Perhaps a better model is provided by adjectives like ‘eligible’ and ‘qualified’, which connote a threshold that varies, depending on what a person is being said to be eligible or qualified for, but this is not a dependence on facts about the speaker or the context of utterance, at least not in any obvious way. At any rate, contextualists have not yet produced, much less defended, a plausible semantic model for ‘know’.

<sup>3</sup> That is, even without being ambiguous and even without being understood differently by different people, each such term can be used literally to express any of a range of properties, depending on which standard of application is operative in the context of utterance. It is not that there are different standards of wealth, tallness, or flatness but different standards for applying ‘rich’, ‘tall’, and ‘flat’. For each term, the particular property thus expressed corresponds to a range on the same scale, and the operative standard determines which range that is.

<sup>4</sup> This difference between ‘know’ and other clearly context-sensitive terms was first pointed out by Stephen Schiffer. As he explains, according to contextualism 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 contextualism imputes to people (those not privy to contextualism) a certain semantic blindness about ‘know’. Somehow this semantic blindness does not impede their understanding of ‘know’, even though they do not notice shifts in its use, even when they themselves use it more or less stringently.

<sup>5</sup> Here I am going along with the pretense that contextual variability is context dependence, as if it is context itself that determines the semantic value of a contextually variable expression. In my view, however, it is more accurate to say that the speaker’s intention is what determines it and a mistake to count the intention as an element of the context. So, strictly speaking, ‘contextualism’ is a misnomer. For more on the uses and abuses of context, see “Context *ex machina*” (Bach 2005).

<sup>6</sup> Contextualists differ as to what shifts: strength of justification, extent of relevant alternatives, or range of possible worlds in which the truth is tracked (this difference is investigated by Jonathan Schaffer (forthcoming), who argues what shifts are relevant alternatives). I’ll ignore this difference here. Also, it is not obvious that standards, however conceived, form a linear ordering, but assuming that they do simplifies the discussion, as when contextualists speak of epistemic standards being raised or lowered.

<sup>7</sup> Since what varies is the standard E, the content of ‘knows’ as it occurs in the locution ‘knows-by-E’ is fixed. The idea is that once the locus of contextual variation is identified and the

appropriate relativization is made explicit, there is no further need to treat the content of ‘knows’ itself as contextually variable.

<sup>8</sup> Mark Heller is a clear exception. As he 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).

<sup>9</sup> The strategy discussed here is developed most fully and clearly by Keith DeRose (1995), and the illustrative argument used here is essentially the same as his.

<sup>10</sup> Here we are contrasting skeptical contexts with ordinary, Moorean ones, but the same point applies to more demanding but non-skeptical contexts, in which it is appropriate to raise realistic, as opposed to far-fetched, skeptical possibilities. We will consider those later.

<sup>11</sup> The first premise is still true, although its content is not the same as in the skeptical context, since ‘knows’ now expresses a weaker relation. This assumes a principle of epistemic closure, that knowing that p entails knowing what p entails, or at least what p is known to entail. Closure is accepted by most contextualists, Mark Heller (1999b) excepted.

<sup>12</sup> Where the BBIV possibility is not considered, the proposition expressed by ‘Moore didn’t know that he wasn’t a BBIV’ is false. However, in an ordinary context one cannot use that sentence without transforming the context into a skeptical one, in which case one is not expressing the intended proposition, that Moore didn’t know-by-ordinary-standards that he wasn’t a BBIV. Even so, that proposition, though not expressible by that sentence in a context that is and remains ordinary, is still false.

<sup>13</sup> A number of philosophers, including Richard Feldman (1999 and 2001), Peter Klein (2000), and Ernest Sosa (2000), as well as Kornblith, have registered this complaint.

<sup>14</sup> Here I am ignoring the Putnamesque complication that the contents of our ordinary knowledge attributions would be different.

<sup>15</sup> 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 may be too low and that the subject does not know even without attributor raising the standards. If you think that, then add further stipulations on the low-standards version of the case or else pick a different example, such as one involving ordinary perceptual identification or recollection of a simple fact.

<sup>16</sup> In my view there is a general reason not to rely too heavily on seemingly semantic intuitions, namely that they are often responsive to pragmatic factors (Bach 2001), but that is not the problem in the present case.

<sup>17</sup> It is interesting to note that when introducing skeptical invariantism, Peter 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 skeptic’s ultra-high standards so much as the strength of the doxastic condition on knowledge, arguing that it requires “absence of doubt or doubtfulness.” The idea, I take it, is that the belief be settled (this assumes that beliefs can be unsettled).

<sup>18</sup> Of course, the very fact of its being raised, or its coming to mind and staying there, may make it seem plausible, especially if one doesn’t know how to eliminate it. For more on this idea, see section 5 and also Hawthorne (2004: 170-172).

<sup>19</sup> 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 warranted assertibility arguments or WAMs, as DeRose (2002) calls them. But these are sophisticated WAMs, not lame ones of the sort that DeRose thinks moderate invariantism is stuck with. My defense of moderate invariantism in the face of contextualist examples does not rely on WAMs at all.

<sup>20</sup> Nichols, Stich, and Weinberg (2003) offer strong evidence that epistemic intuitions are not nearly as universal or robust as contextualists dogmatically assume.

<sup>21</sup> 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.

<sup>22</sup> The difference between contextualism and SSI is relevant to a new argument for contextualism recently put forward by Keith DeRose (2002), based on Tim Williamson's knowledge account of assertion (2000: ch. 11). As DeRose sums up this 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)

Evidently, for 'P' to be warrantably assertible by someone is for that person to know that P, since that is the truth-condition for one's utterance of 'I know that P'. Unfortunately, this "powerful argument for contextualism" applies only to first-person cases, in which attributor and subject are the same (as in DeRose's examples). In such a case, obviously, what makes 'P' warrantably assertible by the attributor, that the *attributor* knows that P, is not the truth-condition for his utterance of 'S knows that P', since S is someone else. So this argument seems not to discriminate between contextualism and SSI. For that reason, I won't consider it further.

Also, although I cannot go into detail on this here, it seems unnecessary to posit what Williamson calls a "knowledge rule" on assertion. It seems to me that the only relevant rule on assertion is belief, since an assertion essentially is the expression of a belief; there is a separate knowledge rule or, rather, norm on belief itself. So the knowledge rule has no independent status — it's the relative product of the belief rule on assertion and the knowledge norm on belief.

<sup>23</sup> Jeremy Fantl and Matthew McGrath (2002), John Hawthorne (2004), and Jason Stanley (forthcoming) have argued for this claim in rather different ways and defend different versions of SSI (Hawthorne calls it "sensitive moderate invariantism"). Although these differences are interesting and important, they will not be discussed here, for the point I wish to make is independent of these differences. However, one difference should be noted. In Fantl and McGrath's view and in Hawthorne's (insofar as he actually endorses it), it is what the subject thinks is at stake that matters, whereas in Stanley's view, what matters is what the stakes actually are (even if the subject is unaware of them). The latter view implies, contrary to what I assume in



the text, that there need be no difference in doxastic threshold between the two people being compared. However, I just don't share Stanley's intuition about the relevant cases. It seems to me that one can know something even if, unbeknownst to one, there is good practical reason, because of the cost of being wrong, to make sure. It seems coherent to suppose that can person know something even if he would be well advised to make sure of it.

<sup>24</sup> I am following common usage in contrasting 'practical' with 'epistemic' reasons, but obviously it is misleading, and indeed prejudicial against SSI, to use 'epistemic' here. After all, if, as SSI says, practical reasons can affect whether or not one is in a position to know, they *would* be epistemic (or at least could be).

<sup>25</sup> Sufficient on whatever the correct account of knowledge might be. This may include both evidence one possesses and, to allow for Gettier cases, evidence one does not possess.

<sup>26</sup> Hawthorne observes that "for pretty much any proposition of which we are convinced, we will be inclined to accept a bet against it given the right odds and reckon ourselves perfectly rational in doing so" (2004: 176n), but he does not pursue the possibility that "the sketched connection [based on the assumption underlying SSI] between knowledge and practical reasoning is only roughly correct."

<sup>27</sup> Hawthorne adds, "From this perspective, the only mistake made by the simple moderate invariantist is to suppose that belief of the suitable type is invariably present in the puzzle cases we have been considering" (2004: 173), but *we* have not been supposing this.

<sup>28</sup> According to Lewis, 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 this conception of elimination, it is the experience, not the person having it, that eliminates a possibility.

<sup>29</sup> This is a kind of default reasoning, whereby we jump safely to conclusions without having to verify the many implicit assumptions that we make, so long as we can rely on our ability to think of those and only those possibilities that are worth considering. I have previously defended this conception of default reasoning (Bach 1984) and used it to defend a form of reliabilism about justified belief (Bach 1985).

<sup>30</sup> It is not part of this suggestion that the conditional proposition that is known be explicitly believed. Rather, the idea is that the normality condition is implicitly assumed.

<sup>31</sup> The 'etc.' here goes proxy for other improbable but far-fetched ways in which the normal course of events can be broken, but it does not include skeptical and other fantastic possibilities, such as your car exploding spontaneously or being hoisted by a deranged crane operator on a nearby construction project.

<sup>32</sup> One suggestion along these lines, and consistent with the idea of the previous section, is Harman and Sherman's (2004) proposal, that one can know something without knowing the assumptions on which it rests, provided those assumptions are true and one justifiably takes them for granted.

<sup>33</sup> This idea comports with any fallibilist theory of knowledge which allows that one can know that p even if there are realistic situations in which someone believing the same thing on the basis of the same evidence would not know that p. In most cases, the belief that p, however well justified, would be false. The rare cases in which it is still true that p are Gettier situations. Such cases are possible on a conception of knowledge that is not only fallibilist but externalist as well.

<sup>34</sup> As I have pointed out in other connections (Bach 2001 and 2005), a great many of our everyday utterances are to be taken as not fully explicit, that is, as if they contained words or phrases that they do not actually contain. Even when all its constituent expressions are used literally, such an utterance as a whole is not literal, although intuitively it may seem to be (see Bach 2002).

<sup>35</sup> Lotteries raises these questions: how can you know that your ticket won't win when you can't know of each of the other tickets that it won't win; and, if you can't know that your ticket won't win, how you can you know, as it may seem that you can, that you won't get rich soon? Personally, I find the straight lottery cases unproblematic: you don't know that your ticket won't win, and you don't know that you won't get rich soon.

<sup>36</sup> Note that this is not an alternative to your car being in your driveway, since your car could have been taken on a joy ride and returned to your driveway. But in that case you would be in a Gettier situation and wouldn't know that your car is in your driveway.

<sup>37</sup> Many thanks to Tim Black, Jessica Brown, Mark Heller, David Hunter, Ray Elugardo, John MacFarlane, Patrick Rysiew, and Jonathan Schaffer for very helpful comments, not all of which I could take into account, on a predecessor of this paper, "The Emperor's New 'Knows'." Tim commented on that paper at the 2004 Inland Northwest Philosophy Conference, and Mark and Jonathan commented on it at the 2004 Bellingham Summer Philosophy Conference.