

Quantifiers and Epistemic Contextualism

Jonathan Ichikawa
ichikawa@gmail.com

Forthcoming in *Philosophical Studies*

Penultimate Draft: 25 May, 2010. Please refer to published version.

First submitted 13 March, 2009

Abstract. I defend a neo-Lewisian form of contextualism about knowledge attributions. Understanding the context-sensitivity of knowledge attributions in terms of the context-sensitivity of universal quantifiers provides an appealing approach to knowledge. Among the virtues of this approach are solutions to the skeptical paradox and the Gettier problem. I respond to influential objections to Lewis's account.

David Lewis (1996) articulated a contextualist approach to knowledge on the model of context-sensitive quantifier domain restrictions. Among the virtues claimed for this approach are solutions to the skeptical paradox and the Gettier problem. I believe that the central insight of Lewis's project was importantly correct; my present project is to defend a neo-Lewisian contextualism, retaining Lewis's emphasis quantifier domain restriction. 'Knows' is a modal term, requiring a certain condition to hold among all the possibilities contained in its modal base; which possibilities are included is in part a context-sensitive matter.

1. Context-Sensitive Quantifier Domain Restriction

The crowd cheers as a clown pops out of the small car. 'Wow,' says Nephew, 'there was a clown in that car!' Aunt has seen this sort of thing before: 'Keep watching—not all the clowns are out yet.' Sure enough, out pops a second clown, and then a third, fourth, and fifth. 'I can't believe so many clowns fit inside! Are there any more clowns?' asks Nephew. 'No, all the clowns are out now,' says Aunt. Sure enough, no more clowns jump out of the car. The driver honks the horn and leaves the ring. Aunt spoke truly; all the clowns were out.

Meanwhile, on the other side of the noisy ring, Daughter and Father are watching the same clown car. Daughter cares little for clown tricks, but she loves cars, and often fantasizes about driving them. Her eyes are fixed on Drivo, the clown driving the car. After the fifth clown jumps out, Father notices that Daughter is still watching the car: 'all the clowns are out now,' he says, attempting to redirect her attention to the juggling clowns. 'No,' says Daughter, 'not all the clowns are out. See, there is still a clown driving the car.' Daughter speaks truly; there *is* still a clown in the car—Drivo—so not all the clowns are out.¹

¹ What of Father's utterance? Father's intention was similar to Aunt's—he didn't mean to be talking about the driver. So perhaps we should consider his sentence true. But Daughter, whom we've already agreed spoke truly, also took herself to be *disagreeing* with Father, as indicated by her use of the word 'no'. How to treat such apparent disagreement is something of a vexed issue—its parallel in the case of knowledge is much celebrated, but this case shows that the phenomenon also occurs in less controversial cases of context-sensitive discourse.

We seem to have the beginnings of a puzzle. Aunt truly says 'all the clowns are out', while Daughter truly says 'not all the clowns are out'. And Daughter and Aunt are talking about the very same situation at the very same time. But this puzzle is not too puzzling. It is widely recognized that English quantifiers such as 'all' contain context-sensitive domain restrictions. Daughter and Aunt were using 'all' to range over distinct domains. In particular, Daughter's domain, but not Aunt's, included the driver of the car. Whether a clown counts as part of 'all' the clowns depends in part on features of the speaker's conversational context; contextualism about 'all' is true.

Of course, unknown facts about the clowns can also influence whether a clown counts as part of 'all' (in some speaker's mouth) the clowns. Consider a world w_1 , in which Aunt and Nephew are in just the same situation with just the same intentions as the actual w_0 , watching a small clown car, but where the act was meant to involve not five but six clowns. But in w_1 , the sixth clown, Stucko, got stuck inside the car. The audience could not see that there was another clown trying to get out, and the driver eventually drove off, aborting Stucko's appearance in the ring. In this world, when Aunt says 'all the clowns are out now', she speaks falsely. Stucko is a counterexample, even though, as before, Drivo isn't. Stucko is part of Aunt's domain in w_1 , but outside it in w_0 , even though, from the point of view of Aunt and her conversational participants, everything appears the same. So the domain restriction for English quantifiers depends not only on the conversational context of the speaker; it depends also on facts about the situation—including facts of which the speaker may be unaware.

What I have said thus far ought not, I think, to be controversial. The contextualist treatment of quantifiers is well-established, and the verdicts about the three cases offered seem straightforward.² Of course, there are more questions unanswered: how do conversational contexts interact with speaker intentions and features of the world to produce the extensions of the relevant domains? For instance, what *makes it the case* that Aunt's domain in w_1 includes Stucko? That's difficult to say with any precision. A tempting line suggests that it might have something to do with a kind of salient similarity between Stucko and the other clowns to whom Aunt explicitly attended. This seems to be the difference between Stucko's situation in w_1 (where he, like the other five clowns, is hiding in the clown car), where he is included in the domain, and w_0 (where, let us suppose, he was sleeping at home), where he isn't. Likewise, what *makes it the case* that Daughter's domain includes Drivo, while Aunt's (in both w_0 and w_1) doesn't? Plausibly, something having to do with Daughter's attention: she, unlike Aunt, is watching Drivo and is interested in him. An attempt at a rigorous articulation of the rules that dictate how to determine whether a clown counts, for the purposes of a given utterance, as one of 'all' the clowns would represent an ambitious metasemantic project; I will not presume to undertake it here. Failure to complete this project ought not to undermine the plausible sketch of a contextualist semantics for 'all' gestured at above.

Contextualism about quantifiers is widely accepted; contextualism about 'knowledge' is controversial. David Lewis offered a contextualist account of knowledge in terms of contextualist quantifiers. I find Lewis's approach appealing; however, Lewis undertakes

² This is not to say that contextualism about quantifier domains has no detractors; of particular note is the 'Semantic Minimalism' of Cappelen & Lepore (2005). Defending contextualism about quantifiers is beyond my present scope (but see Stanley & Szabó (2003)). My present conclusion may be thought of as the conditional claim that *if* contextualism about quantifiers is well-motivated, then so is contextualism about knowledge.

particular commitments that have made his account easy for many philosophers to reject.

My present project is to rehabilitate a Lewisian contextualist approach to knowledge. If we keep sight of the examples of context-sensitive quantifier domains mentioned above, we will find that many objections to Lewis are either misplaced or aimed at inessential features of Lewis's account.

2. Lewis on Knowledge

My brand of contextualism is inspired by and related to David Lewis's. According to Lewis (1996), knowledge attributions work like universal generalizations. He gives us this account of knowledge:

S knows proposition P iff P holds in every possibility left uneliminated by S's evidence. (551)

Lewis's 'knowledge' inherits its context-sensitivity from the context-sensitive 'every possibility'. On Lewis's view, there is a class of possibilities that grows and shrinks according to conversational context; our knowledge attributions are true when the subject's evidence eliminates all of the members of that class in which the object of knowledge is false. Lewis goes on to offer a series of rules to determine which possibilities count as part of 'every' possibility: The *Rule of Actuality* demands that the actual world is always in the domain; this entails the factivity of knowledge. The *Rule of Belief* has it that a subject's belief world is in the domain; we never properly ignore that which the subject believes to obtain. The *Rule of Resemblance* has it that worlds sufficiently similar, in a salient way, to a world in the domain are also in the domain. Finally, the *Rule of Attention* has it that any cases we are considering as possibilities are in the domain. It is primarily this rule that gives Lewis's account its contextualist feature—for what cases are treated as possibilities depends partly on the conversational context. Often, we treat cases as relevant possibilities only once they've been mentioned. (The Rule of Resemblance also has contextualist implications; I shall return to this rule in detail in discussing an objection from Stewart Cohen in §8 below.)

One immediate problem for Lewis is that his criterion for knowledge that *p* is met any time the evidence entails that *p*, relative to the relevant set of possibilities. So Lewis's account entails that any time a subject has evidence E, she knows that she has evidence E. Proof: Suppose that S has E. There are no possibilities in the relevant context set (or, in this case, even beyond) in which she has the evidence she actually has, and she also fails to have E. This follows straightforwardly from the fact that, for any evidence E, that the subject has E entails that E. So our subject, having E, thereby knows she has E. This result is implausible.

Another odd implication of Lewis's account is that knowledge does not entail belief. Bafflingly, Lewis seems happy explicitly to embrace this result: 'I even allow knowledge without belief, as in the case of the timid student who knows the answer but has no confidence that he has it right, and so does not believe what he knows.' (556) Lewis's remarkable ability to embrace counterintuitive consequences notwithstanding, we should not accept knowledge without belief.

I therefore propose the following modification of Lewis:

S knows that *p* just in case, for some evidence E, (i) S believes that *p* on the basis of E, and (ii) all the E cases are *p* cases.

This is the view I will defend, although it will require clarification below. My proposal is very much in Lewis's spirit, and retains the context-sensitivity in the quantifier in my condition (ii). It adds an invocation of basing—something it would be surprising to see the correct theory of knowledge do without.

3. Lewis's Rules

My proposal so far is a modification of Lewis's; it offers an account of knowledge in terms of a context-sensitive 'all cases'. It is uncontroversial that the semantic value of a sentence of the form 'S's evidence eliminates all the not-*p* cases' depends in part on its conversational context. English quantifiers like 'all' are context-sensitive. Lewis's substantive claim is that knowledge attributions vary in semantic value the same way that universal generalizations about 'all cases' do.

Lewis goes on to attempt to articulate how the conversational context, together with the subject's circumstances, determines the relevant modal base. He posits a system of rules, alluded to above: the rules of actuality, attention, belief, and resemblance place constraints on which possibilities must be included; the rules of reliability, method, and conservatism articulate defeasible conditions that may exclude possibilities. These particular rules have proven at best controversial.

The framework against which I suggest we consider all of these challenges is this. Lewis offers, in 'Elusive Knowledge', two independent views: (a), a modal account of knowledge, in terms of the context-sensitive 'all cases', and (b), an account of the context-sensitivity of the latter, for which he posits a system of rules. In my view, Lewis's (a) represents a well-motivated insight, and something like it is true; by contrast, his (b) represents an extremely ambitious metasemantic project whose execution may be deeply flawed. Although I think that Lewis's rules gesture at plausible features of the dynamics of conversation, it is doubtful that any system so simple could be better than rather approximately correct.

Compare the case with sentences about 'all the clowns', like those in the opening of this paper. It is one thing to describe their quantifiers as having context-sensitive restricted domains; it is another to articulate the rules that determine which clowns count. Lewis's project was, I think, unnecessarily ambitious; the controversial particulars have distracted many philosophers from the central insight.

Consider a parallel in an evaluation of the sentence, 'Papageno finished the wine.' Plausibly, it is context-sensitive in a way similar to Lewis's knowledge: it is true just in case, roughly, Papageno drank some wine such that, after he drank that wine, *all* of the wine had been drunk. That's part (a).

With sufficient ambition, we could try to fill in part (b) of the story as well, articulating a theory of how the circumstances of evaluation, and the facts about Papageno, determine which wine counts for the purpose of whether he drank 'all' the wine. Rules corresponding to Lewis's rules might appear attractive: any wine that we're thinking about counts; any wine that is obviously available to Papageno counts; any wine whose availability is relevantly similar to wine that is obvious to Papageno counts, etc.

Any list of rules that captures even most of the data will be rather baroque and invite charges of ad hocery. Some of the rules will cite features of the speaker (which wine is salient in the conversation), while some will cite facts about Papageno of which the speaker might be

ignorant (which wine Papageno intends to drink tonight). This is true of quantifiers generally, as the opening clown cases make clear.

This shows that the correlate of Lewis’s project (b) is complicated and difficult to formalize in terms of simple rules—not that the correlate of (a) is misguided. An objection to Lewis’s part (b) does not automatically count against part (a). I shall return to this refrain several times in what follows.

4. Attending and Ignoring

Lewis says that the mere mention of a skeptical possibility thereby *forces* it into the domain; that the mere act of considering possibilities destroys knowledge.³ (Hence, ‘elusive’.) He is forced to describe his own project, then, as an attempt at ‘saying what cannot be said’ (1999, 566). This treatment has a parallel possible view about ‘all the wine’: any wine that has been mentioned or thought about in the context of the speaker is forced into the relevant domain, even if the speaker doesn’t want it there, and it could otherwise have been legitimately ignored.

This result is implausible. The mere mention of some wine does not force it into the domain. Suppose that you and I are at dinner, discussing Papageno (who is far away from us). We see that his glass is empty; you suggest, ‘maybe Papageno will pour another glass now.’ I reply, ‘no, Papageno has already finished the wine.’ Papageno’s last bottle is empty, so I speak truly even if, as we speak, I am opening a new bottle at our table and pouring you a fresh glass. Our wine—to which we are attending—is outside this domain. Plausibly, this is true of possibilities and knowledge, too; mere mention of a possibility isn’t enough to make it relevant.⁴

The point I wish to emphasize once again is merely that the particular machinery that explains how the context-sensitive quantifier’s domain depends on context will be complicated, and there is room, consistent with the account of knowledge I defend, to disagree on these sorts of particulars. It is also worth remembering that the particular challenges of articulating which cases count as part of ‘all’ the not-*p* cases are general challenges that extend well beyond difficulties with Lewisian contextualism about knowledge; it’s also difficult to articulate which wine counts as part of ‘all’ the wine.

5. Evidence

Lewis’s account of knowledge invoked evidence; my modification of Lewis does too. It is controversial which propositions are part of one’s evidence; Lewis clarifies that his evidence is internal—perceptual experience, apparent memory, and the like. He is sanguine as to just what belongs on the list—‘If you want to include other alleged forms of basic evidence, such as the evidence of our extrasensory faculties, or an innate disposition to believe in God, be my guest’—but it seems clear that he thinks, whatever counts as evidence, it must be at least

³ This, of course, is a significant mischaracterization of Lewis own view; to force a context in which ‘knows’ does not apply is no more to destroying knowledge than is to force a context in which ‘everybody’ applies is to destroy everybody.

⁴ Compare Hawthorne 2004, p. 64: ‘We don’t want to say that people watching [*The Matrix*] are automatically in a context where they cannot truly say ‘I know I’m in a movie theatre.’ See also Oakley (2001).

introspectively recognizable. (553) In defending a cousin of Lewis’s view, I don’t mean to commit to his internalist conception of evidence.

One way—not a particularly interesting one—to have my view about knowledge along with externalism about evidence is to understand ‘evidence’ in the account of knowledge as a term of art. We should invoke not all our evidence, but all our evidence*—evidence* is what Lewis thought evidence was. We could, and probably should, write out the misleading word, thus: *S knows that p just in case, for some set of perceptual and/or memory experiences {e₁, ..., e_n}, (i) S believes that p on the basis of {e₁, ..., e_n}, and (ii) all the cases where {e₁, ..., e_n} obtain are p cases.*

Another, more interesting way to reconcile Lewis-style contextualism with evidence externalism is simply to allow the external evidence to play the role given to ‘evidence’ in the characterization of knowledge as stated.

We might worry about circularity or triviality results if we are open to certain externalist accounts of evidence. Consider Timothy Williamson’s—your evidence is all and only that which you know.⁵ Or, as the contextualist would put it, “S knows *p*” is true in a context just in case “*p* is part of S’s evidence” is true in that context. Can we plausibly adapt even a knowledge-based account of evidence to this account of knowledge? In a word: yes. Two observations mitigate against any counterintuitiveness of this response. First, I am not intending to offer an *analysis* of the concept knowledge, breaking it into components that are conceptually prior to knowledge. This project strikes me as both hopeless and uninteresting. Second, my introduction of a basing requirement makes the account much friendlier to a knowledge-first treatment of evidence. Here again is Lewis’s own original formulation:

S knows that *p* just in case S’s evidence eliminates all the not-*p* cases.

Given Williamson’s account of evidence, Lewis’s account becomes all but vacuous:

S knows that *p* just in case S’s knowledge entails that *p*.

(This isn’t *quite* vacuous, because it rules out the case where you fail to know that *p* because you fail to recognize the truth that *p* is entailed by your knowledge; of course, this is the sort of knowledge that I objected to in §2.) However, the invocation of a basing requirement changes the shape of the view. On my account,

S knows that *p* just in case, for some evidence E, (i) S believes that *p* on the basis of E, and (ii) all the E cases are *p* cases.

If we plug in E=K, we get:

S knows that *p* just in case, for some known propositions {e₁, ..., e_n}, (i) S believes that *p* on the basis of {e₁, ..., e_n}, and (ii) all the cases where {e₁, ..., e_n} are true are *p* cases.

This principle is not vacuous in the way the previous one was. It places a genuine constraint on knowledge: in order to know that *p*, you must believe that *p* on the basis of some sufficient

⁵ Williamson (2002), ch. 9.

knowledge. This framework leaves open the possibility that there is basic knowledge; this would be the limiting case in which the evidence is known on the basis of itself. This invites a foundationalist framework of basic knowledge: which propositions are legitimately known in this way?⁶

On this treatment, ‘evidence’ will be context-sensitive, too.⁷ This is no vicious circularity. My knowledge depends on whether that which I may legitimately rely upon is strong enough to eliminate all the relevant alternatives; the conversational context of the knowledge attribution affects not only which alternatives are relevant, but also that which I may, for such purposes, legitimately rely upon.

So the neo-Lewisian approach to knowledge that I favor is compatible with a wide range of views about evidence, including E=K.

6. Possibilities

S knows that *p* only when his evidence eliminates all the not-*p* cases. What is a case? For Lewis, it’s something like a possible world, or a set of possible worlds. An odd implication, then, is that there are no distinct impossible cases; it is therefore very easy to know the proposition that *Hesperus is Phosphorus*; this proposition is identical to the proposition that *water is H₂O*, and that *triangles have three sides*. This is a problem common to philosophical views that individuate contents coarsely. Lewis’s (1999) response, which is a standard one, is:

[T]he necessary proposition is known always and everywhere. Yet this known proposition may go unrecognised when presented in impenetrable linguistic disguise, say as the proposition that every even number is the sum of two primes. (552)

This course-grained treatment of propositions is also optional on the present view. We may well take Lewis’s approach, or we may choose instead to invoke impossible worlds, or structured propositions, where some propositions are impossible. I am inclined towards the latter options—we can thereby preserve the intuition that Lois doesn’t know that Clark is Superman, and that the possibility that Clark is not Superman is uneliminated by her evidence. But I mean to remain officially neutral on this point, too.

The point of the preceding sections has been to motivate that many of the controversial elements to Lewis’s account of knowledge are inessential to the core contextualist insight, which I endorse.

7. Solving the Skeptical Paradox

This contextualist treatment of knowledge resolves a familiar form of a skeptical paradox:

⁶ Other options than foundationalism are also available; one may respond to the classical regress challenge in any of the usual ways. One might, for instance, follow Peter Klein (2007) and demand an infinite chain of known reasons, or admit circular chains of evidence on the model of coherentism.

⁷ At least it does so if the statement of E=K is meant to hold true in all contexts; other options are also available. For example, one could, if one liked, identify evidence with what is ‘known’ relative to some particular context; the resulting view is also consistent with my approach to knowledge.

- (1) I know that I have hands.
- (2) I don’t know that I’m not a handless brain in a vat.
- (3) It is absurd that I know that I have hands but don’t know that I’m not a handless brain in a vat.

Each premise is plausible, but their conjunction appears contradictory. The contextualist account allows each premise consistently to be true; an utterance of (1) expresses a truth when made in normal, non-skeptical contexts, because cases in which the subject is massively deceived are not part of the domain. (There are cases in which he has no hands, but these are cases his evidence eliminates.) An utterance of (2), however, brings attention to the skeptical hypothesis; by the Rule of Attention (or a suitably weakened version of it—see §4 above), the domain expands to include that hypothesis, and the knowledge claim is false: there is a case in the domain, uneliminated by the subject’s evidence, where he has no hands. Likewise for (3); in any context, it is false—indeed, absurd—to utter both (1) and (2) together.

So the paradox is diffused by assimilating it to this non-paradox:

- (1’) I drank all the beer.
- (2’) I didn’t drink the beer we left in the store.
- (3’) If I drank all the beer, then I drank the beer we left in the store.⁸

Jason Stanley (2005) offers a critique of various forms of contextualism about knowledge, including the charge that this Lewisian contextualist line cannot resolve the paradox in the way I (and Lewis) have suggested. Quantifiers, Stanley says, can shift domains in and out in the middle of conversational contexts; we are happy to accommodate speakers who change the scopes of their quantifiers midway through conversations or even sentences. So, for instance, there is nothing problematic with:

Every sailor waved to every sailor. (60)

(In a context in which two ships of sailors have just passed one another.) We accommodate, treating each ‘every’ as having a different scope, and so not requiring that every sailor waved to himself, or to the other sailors on his own ship.

Stanley also gives us this dialogue to illustrate how easily quantifiers can shift mid-conversation:

- A. Every van Gogh painting is in the Dutch National Museum.
- B. That’s a change; when I visited last year, I saw every van Gogh painting, and some of them were definitely missing. (65)

⁸ Note also that this treatment explains some dynamic features of knowledge attributions; to claim (1) or (1’) is, in many contexts, felicitous, and (2) and (2’) may well follow in the conversation. But statements of (1) or (1’) are much more problematic when given *after* (2) and (2’). This apparent dynamic feature is further evidence that there is a contextualist element at work in the case of knowledge.

There are three quantifiers: the first and third range over all the van Gogh paintings in existence; the second ranges only over those that were in the museum last year.

Since quantifiers shift so readily, Stanley says, Lewis’s view should predict that knowledge attributions should also shift very easily. So according to Stanley, the view should predict that there is nothing wrong with ‘abominable conjunctions’⁹ like:

S knows that he has hands but S does not know that he is not a handless brain in a vat.

If knowledge attributions are just quantifiers, Stanley says, then why don’t we adjust the understood domain restrictions to make sentences like this one unproblematic? The objection to contextualism, then, is that it wrongly predicts that such abominable conjunctions express truths, and therefore fails to resolve the skeptical paradox.

But this is unfair to the contextualist. Stanley has proven that sometimes we accommodate assertions by shifting domains; that doesn’t mean that we have unlimited flexibility in this respect. It’s not the case that our domains shift to *whatever* makes our utterances true in all cases. And indeed, in the apparently-relevant cases, domain shifting does not render the natural reading true. Take Lewis’s gloss on the abominable conjunction above:

All of S’s uneliminated possibilities are hand-possibilities, but some of S’s uneliminated possibilities are handless-brain-in-a-vat possibilities.

This conjunction is just as abominable as the knowledge one. It’s a bit of a mouthful, so consider a simpler conjunction with the same form:

All of the bottles are on the table, but some of the bottles are in the fridge.

Because these quantifier-conjunctions are abominable, Lewis’s theory does not predict that the relevant knowledge-conjunctions are unabominable. It is no embarrassment, then, that the knowledge-quantifiers don’t shift in a way to make abominable conjunctions true.

8. The Rule of Resemblance: Gettier Cases and Lotteries

According to Lewis, ‘knows’ is a context-sensitive modal whose base is fixed by a series of rules. ‘S knows p’ is true in a context just in case S’s evidence entails p, relative to the modal base, which depends in part on context. The rules of actuality and belief are subject-sensitive; they ensure that any possibility that does obtain in the subject’s situation, or that the subject believes to obtain, or ought to believe to obtain, are among the relevant worlds. The rule of attention is ascriber-sensitive; it generates context-sensitivity for knows. Possibilities to which the speaker is attending are included in the modal base.

The Rule of Resemblance is a bit trickier. Possibilities that saliently resemble other relevant possibilities are included into relevance. Similarity facts depend and supervene on the intrinsic qualities of the possibilities, so that element of the rule is subject-sensitive and ascriber-insensitive. However, it’s not brute, objective resemblance facts that Lewis cites: it’s *salient* resemblance. Salience, presumably, is going to be an attributor-sensitive property. So the rule of resemblance is going to introduce some context-sensitivity too. Which possibilities

⁹ The term is due to DeRose (1995), pp. 27-28.

are *saliently* similar to, e.g., actuality, depends in part on the speaker’s context.

This sensitivity corresponds to certain features of the lottery paradox. Although we standardly judge that we do not ‘know’ that a given lottery ticket will lose, there are some contexts in which we are happy to claim ‘knowledge’ of propositions that transparently entail that it will. Lewis assimilates this shift to a shift in the salience of resemblances:

Pity poor Bill! He squanders all his spare cash on the pokies, the races, and the lottery. He will be a wage slave all his days. We know he will never be rich. ... Salience, as well as ignoring, may vary between contexts. Before, when I was explaining how the Rule of Resemblance applied to lotteries, I saw to it that the resemblance between the many possibilities associated with the many tickets was sufficiently salient. But this time, when we were busy pitying poor Bill for his habits and not for his luck, the resemblance of the many possibilities was not so salient. (565)

Lewis also uses the Rule of Resemblance to diagnose Gettier cases. Why do Gettier subjects fail to know p? Because there is a possibility, saliently close to actuality, in which their evidence obtains, but not-p. For example, Henry in fake barn country is in a world in which there are fake barns about; that world saliently resembles the possibility in which he himself is looking at a fake barn, which his evidence does not eliminate; so he does not know he faces a barn.

Stewart Cohen (1998) has objected to this combination of Lewis’s views: the subject-sensitivity of salient resemblance and the application of the Rule of Resemblance to Gettier cases. The ignorance of Gettiered subjects, Cohen says, cannot plausibly be construed as a context-sensitive matter. In no context can ‘S knows p’ be true, where S stands to p in the way articulated in a Gettier-style thought experiment.¹⁰ In particular, Cohen charges that Lewis is unable to offer the correct diagnosis of attributions of knowledge to subjects in cases where a subject is in a Gettier case but the attributor is in no position to appreciate that this is so.

I believe that Lewis’s approach has the resources to answer Cohen’s challenge. However, before developing this claim, let me emphasize again that, when we work with the generalized view I have been defending, this objection does not even get off the ground. In the classification given in §3, Cohen has offered an objection to the peripheral (b), not the central (a). According to my view, the knowledge attribution is true only if the subject’s evidence eliminates ‘all’ the possibilities in which the content of the belief is false, where that ‘all’ is restricted to some class of relevant possibilities fixed by the conversational context. When subjects are in Gettier cases, the relevant possibilities include the nearby skeptical cases (e.g., looking-at-fake-barn cases); there is no commitment to the strong claim that such nearby cases are included if and only if they are similar to actuality in a way salient to the attributor.

¹⁰ Cohen occasionally states the objection in the object language, as here: “Intuitively, because of his [Gettier] situation, S_i fails to know he sees a sheep, regardless of what is salient to the speaker of the context.” (301) Presumably, this is to be interpreted in the metasemantic way suggested; the literal fact given here is not in tension with contextualism of the relevant sort. Compare Lewis’s own remark that his talk of knowledge vanishing with context shifts is an “expository shortcut, to be taken with a grain of salt.” (566)

One might object that without principles explaining why some similarities are sufficient for inclusion, but others are not, the claim that just the ones that falsify Gettier knowledge count is suspiciously *ad hoc*. But it is not *ad hoc*; it is an instance of a general phenomenon of which we're already aware. Recall the opening example of clown cars. Aunt sees five clowns pop out and utters, 'all the clowns are out'; her utterance is false in w_1 , because Stucko, a clown who was meant to pop out, is stuck in the car, and is therefore included in the domain of her quantifier. Had there been no such hidden clown (as in w_0), the domain would have been smaller, even though no facts about Aunt's salience change between w_1 and w_2 . The *fact* that Stucko is in the car trying to get out—like the fact that S is in a Gettier case—results in the inclusion of Stucko into the one domain, and the possibility that there is a rock but no sheep into the other. This even absent salience for the attributor that such conditions obtain. This is not to deny, of course, that facts about Aunt's context also matter; her intention plays a central role in determining which features of the clowns—in this case, that they are meant to jump out of the car—suffice for inclusion in the domain.

So even if Cohen's objection against Lewis were sound, it is not clear that it would tell against my neo-Lewisian view. However, in this case, I believe Lewis himself has the resources to resist the implication that Cohen suggests he must accept.

Cohen writes:

S sees what appears to be a sheep on the hill. But what S actually sees is a rock that looks, from that distance, to be a sheep. It happens though, that behind the rock, out of S's view, is a sheep. ...

Now consider A, standing next to S, who is unaware that S sees only a rock. The resemblance between the possibility that S sees a rock that looks like a sheep and actuality is not salient for A. A is not aware that S is in a Gettier situation of any kind. So according to Lewis's view, in A's context of ascription, the possibility that S sees merely a sheep-shaped rock can be properly ignored. Thus on Lewis's view, A truly ascribes knowledge to S. A can truly say 'S knows there is a sheep on the hill'. (297)

The two possibilities that are relevant here are the actual possibility w_1 —that S sees a rock on a hill on which there is also a sheep—and the skeptical possibility w_2 —that S sees a rock on a sheepless hill. Lewis wants w_2 to count as relevant by virtue of its similarity to w_1 , but Cohen argues that, in the case envisaged, that move is unavailable, since the subject, being ignorant of the relevant features of actuality, is in no position to appreciate the similarity. The assumption, then, appears to be something like this: in order for two possibilities saliently to resemble one another, it must be salient to the attributor that those possibilities saliently resemble one another. But Lewis is not committed to this condition on salient similarity. Indeed, he seems to commit himself to its negation when he writes:

We must apply the Rule of Resemblance with care. Actuality is a possibility uneliminated by the subject's evidence. Any other possibility W that is likewise uneliminated by the subject's evidence *thereby resembles actuality in one salient respect*: namely, in respect of the subject's evidence. That will be so even if W is in other respects very dissimilar to actuality—even if, for instance, it is a possibility in

which the subject is radically deceived by a demon. (556, my emphasis)¹¹

Lewis does not here stipulate anything about an ascriber's attention to particular possibilities; it needn't be salient to anyone that actuality resembles the possibility that I am a brain in a vat; the argument does not depend on anyone's thinking about brains in vats at all. The skeptical possibility is saliently similar to actuality in respect of the subject's evidence—and that respect is salient. The Rule of Resemblance requires salience of some property with respect to which a possibility is similar—not that the skeptical possibility itself be salient, or that the fact that actuality is similar to some skeptical possibility be salient. Cohen writes that the attributor-sensitivity of the Rule of Resemblance "means that features of the context of ascription-facts concerning what resemblances are salient ... will determine which possibilities can not, by this rule, be properly ignored." (295) But this determination claim is much too strong; the context-sensitivity means that salience features of the conversational context are relevant to determining which possibilities can properly be ignored, not that they salience fixes the answer to these questions.

Something like a 'Rule of Similarity' might be at play in governing the domains of quantifiers. Why is Stucko part of Aunt's domain? Plausibly because, *in some relevant respect*, he's similar to the clowns that Aunt intended to be talking about. This is consistent with Aunt's total ignorance of Stucko; it is not salient to Aunt that Stucko is similar to the other clowns. What is salient to Aunt is the relevant property: being a clown that is meant to jump out of this car tonight. Contrast another clown, GemGem, who is the identical twin of the first clown to jump out of the car. She is in *some* respect similar to a car clown, but it is not the relevant one. Notably, this is so regardless of how salient that fact of similarity is. Aunt may have been thinking or even saying, 'that clown who just jumped out looks *exactly* like my friend GemGem'; this puts no pressure whatsoever toward including GemGem in the domain of her quantifier when she says that 'all the clowns are out'.¹²

Salient similarity, then, can be glossed as something like 'similar with respect to the features that are salient'. Facts about the speaker and the conversational context play roles in picking

¹¹ In introducing this complication, Lewis sets out what he recognizes to be a problem for his own view: since there is always some skeptical possibility that is saliently similar to actuality, how is it that any knowledge attributions come out true? Lewis confesses that he does not know how to avoid this consequence without ad hocery. Cohen seems (p. 296) to consider it obvious that the best solution for Lewis is to understand the salient similarity of x and y as requiring that it be salient that x and y are similar. However, as Cohen goes on to show, taking this line brings about further problems for Lewis with respect to Gettier cases. Cohen's moral is that Lewis has serious problems with Gettier cases; mine is that Cohen's interpretation of Lewis is not the best one. As I'll go on to argue, Lewis seems to intend a different notion of salient similarity—one that does not provide an obvious solution to the challenge introduced in the quoted paragraph.

In effect, then, the Lewisian contextualist faces a choice: he can adopt Cohen's suggestion and require a stronger notion of salient similarity, thus avoiding the skeptical challenge here described without ad hocery; or he can use a weaker notion of salient similarity that gets the results Lewis wants about Gettier cases. I side with Lewis in taking the latter option, and therefore inherit his problem of leaving unanswered the question about why 'similarity with respect to evidence' is not usually enough to include skeptical possibilities in the domain. But see fn. 12.

¹² We can also find a problem parallel to the one Lewis admits for himself, discussed in my fn. 11. Every clown in the world resembles the clowns Aunt is focusing on in at least some salient sense—the sense of being a clown. Here, as in the knowledge case, the Rule of Resemblance must be treated with care, and attempts to precisify it will look rather *ad hoc*.

out the relevant similarity relation; thus picked, it has a role to play in determining truth conditions, even if its relata were neither known nor intended by the speaker. So it is that Stucko falsifies Aunt’s claim: he’s relevantly similar to the clowns she explicitly considered (*i.e.*, he’s hiding in the car, preparing to jump out); so likewise is the looking-at-a-fake-barn case relevantly similar to actuality (*i.e.*, he’s in this town, looking at something that looks just like a barn).

So Lewis needn’t be committed, in giving his approach to Gettier cases, to assuming that the fact that actuality is similar to a skeptical case be salient to the attributor of knowledge; what he requires is that the attributor’s context makes salient a particular feature of the subject’s case, where the subject’s situation is similar to a skeptical case in respect of that feature. In Cohen’s rock case, plausibly, that feature is the subject’s reliance on visual evidence in the relevant viewing conditions.¹³

This way of treating the case sits comfortably with Lewis’s own discussion of ‘Poor Bill’. Depending on what features of Bill we focus on, we can be inclined to attribute either knowledge or ignorance of Bill’s future financial situation; if we focus on the evidence to the effect that Bill is playing in the lottery, then we will deny that a subject¹⁴ knows Bill will never be rich, since there are possibilities that are similar with respect to Bill’s gambling in which he wins. But if instead of focusing on the gambling, we focus on Bill’s shiftless habits, we don’t get possibilities in which Bill wins the lottery as saliently similar, and we do properly ignore the case in which he wins.

So contrary to Cohen’s claim, Lewis needn’t be committed to the truth of knowledge attributions to Gettier subjects by attributors who are ignorant of the Gettier situation.

What, however, of Cohen’s more general claim, that ‘S knows p’ should be false for a Gettier subject and proposition in all contexts? Salient similarity *is* context-sensitive; shouldn’t there be possible contexts in which the similarity between the Gettiered actuality and the skeptical case is not salient, even in my weaker sense? While Lewis’s commitments do not, it seems to me, *guarantee* the possibility of a context with such features, it does not seem to offer any straightforward explanation for why there should be none. So if Cohen is right and there are no contexts in which ‘S knows p’ is true of Gettier situations, then Lewis’s version of contextualism fails to explain an important fact.

However, I do not think that we face any great pressure to accept Cohen’s strong claim that there is no possible context in which ‘S knows p’ can be true of Gettier situations. (We have already accommodated, on behalf of Lewis, Cohen’s weaker claim that in the particular cases he discussed with ignorant attributors, ‘S knows p’ is not true of the Gettier situations.) Cohen gives no argument in favor of this stronger claim, offering only that “[s]urely it is very strange to suppose that there is any context of ascription in which one can truly say of S that he knows there is a sheep on the hill. The sentence, ‘S knows there is a sheep on the hill’ looks false (at that world and time), regardless of who happens to be uttering it.” (298)

¹³ Cohen (p. 297, fn. 19) objects to this interpretation of Lewis’s ‘salient similarity’ by arguing that it does not explain why the Rule of Resemblance doesn’t make all skeptical scenarios relevant, resulting in widespread skepticism. But this of course is exactly the situation as Lewis sees it; see previous footnote.

¹⁴ Lewis’s own discussion of Bill is in the first person; he gives the special case where the subject is the ascriber: “we know he will never be rich.”

However, intuitions about Gettier cases of various stripes are much easier to shift than Cohen’s claim here seems to suppose. Lewis emphasized that Lottery intuitions shift when we focus on different features of the evidence—when we’re thinking of the probabilistic nature of the lottery, ‘S knows Bill won’t get rich’ is false, but when we’re thinking of his shiftless, lazy nature, the sentence expresses a truth. It does seem possible to get the same effect with regard to Gettier cases. Suppose someone receives testimony that p from a reliable and knowledgeable informant, but in a room containing many liars. If we think of her as receiving testimony from a knowledgeable source, we’re inclined to say she ‘knows’—if we think of her as receiving testimony from someone in a room containing many liars, we’re inclined to say ‘she fails to know’. Gendler & Hawthorne (2005) offer a catalogue of related ways to shift about knowledge judgments with Gettier cases—they briefly mention the potential applicability of the data discussed to contextualism.¹⁵

Relatedly, some contexts, in which we don’t focus on evidence at all, may deliver truthful ‘knows’ attributions to Gettier situations. Sometimes, for example, we’re not concerned with a subject’s grounds for belief—we’re concerned instead with whether the subject is culpable because knowledgeable. “What did she know, and when did she know it?” Arguably, the answer to this question does not depend on whether she acquired some of her true beliefs in a Gettierized way. If this is right, then there are contexts in which ‘knows’ applies to Gettier situations. We might hold someone responsible for ‘knowing’ that there was a sheep on the hill, even if he were in a situation like the one Cohen described. So the contextualist treatment offers the intuitive result about this case. Of course, the standard response to such sentences is to describe them as loose talk: felicitous, but not literally true. So a theory that does not predict the truth of such sentences, when their subjects are in Gettier cases, does not obviously suffer a great intuitive cost. However, where invariantists must treat such sentences as false but felicitous, contextualists are free to treat them as literally true. (Note that invariantists who are at least moderately skeptical will describe many ordinary knowledge attributions as loose but false, where the contextualist treats them as literally true.)

None of this, of course, comprises anything like a complete story about the relevant dynamic features of discourse about knowledge. Lewis merely gestured at salient similarity; I’ve done little more in this section than slightly to clarify this gesture. More generally, many metasemantic questions about the role of context in picking out a modal base for knowledge attributions remain here unanswered. Some philosophers will be dissatisfied for this reason. In my view, such philosophers are demanding too much at once. To understand knowledge attributions as context-sensitive on the model of quantifiers, and to show that they shift in the same sorts of ways, is a substantive and interesting philosophical conclusion. Of course questions remain unanswered; these questions have unanswered analogues for quantifiers, as I’ve repeatedly emphasized in this paper. Such considerations, therefore, do not undermine the contextualist approach to knowledge, any more than they undermine the widely-accepted contextualist approach to quantifier domains.¹⁶

¹⁵ fn. 20.

¹⁶ For valuable conversations, and for and comments on previous drafts of this project, I am grateful to Jessica Brown, Yuri Cath, Stewart Cohen, Paul Dimmock, Benjamin Jarvis, Carrie Jenkins, Ernest Sosa, Jason Stanley, Brian Weatherston, and Crispin Wright.

References

- Cappelen, H. & Lepre, E. (2005) *Insensitive Semantics*. Malden, MA: Blackwell Press.
- Cohen, S. (1998). Contextualist solutions to epistemological problems: Scepticism, Gettier, and the lottery. *Australasian Journal of Philosophy* 76(2): 289-306.
- DeRose, K. (1995). Solving the skeptical problem. *The Philosophical Review*, 104(1): 1-52.
- Gendler, T. & Hawthorne, J. (2005). The Real Guide to Fake Barns: A Catalogue of Gifts for Your Epistemic Enemies. *Philosophical Studies* 124(3): 331-52.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. New York: Oxford University Press.
- Klein, P. (2007). Human Knowledge and the Infinite Progress of Reasoning. *Philosophical Studies* 134(1): 1-17.
- Lewis, D. K. (1996). Elusive knowledge. *Australasian Journal of Philosophy* 74(4): 549-67. Reprinted in Lewis, D. (1999). *Papers in Metaphysics and Epistemology*, Cambridge: Cambridge University Press: 418-45.
- Oakley, I. T. (2001). A skeptic's reply to Lewisian contextualism, *Canadian Journal of Philosophy*, 31(3): 309-332.
- Stanley, J. (2005). *Knowledge and practical interests*. New York: Oxford University Press.
- Stanley, J. & Szabó, Z. (2003). On Quantifier Domain Restriction. *Mind & Language* 15(2): 219-61.
- Williamson, T. (2002). *Knowledge and its Limits*. New York: Oxford University Press.