

Quantifiers, Knowledge, and Counterfactuals

Jonathan Ichikawa
Arché Philosophical Research Centre
The University of St. Andrews

ichikawa@gmail.com

Forthcoming in *Philosophy and Phenomenological Research*.

PENULTIMATE DRAFT; PLEASE REFER TO PUBLISHED VERSION.

Abstract. Many of the motivations in favor of contextualism about knowledge apply also to a contextualist approach to counterfactuals. I motivate and articulate such an approach, in terms of the context-sensitive ‘all cases’, in the spirit of David Lewis’s contextualist view about knowledge. The resulting view explains intuitive data, resolves a puzzle parallel to the skeptical paradox, and renders safety and sensitivity, construed as counterfactuals, necessary conditions on knowledge.

Introduction

Contextualism about knowledge is motivated by, among other considerations, skeptical paradoxes. Moralès knows that José is not on duty, and sincerely tells Micaëla so. She believes him. It seems that (1) expresses a truth:

1. Micaëla knows that José is not on duty.

Micaëla had never met Moralès before; she accepts his testimony rather unreflectively. We are therefore inclined to judge that (2) is false:

2. Micaëla knows that Moralès isn’t lying about José.

Juxtaposing our verdicts about (1) and (2), we have a puzzle; for their apparent conjunction, (3) is, under the circumstances, abominable¹:

3. Micaëla knows that José is not on duty, but not that Moralès isn’t lying.

Contextualism about ‘knows’ may be thought to provide an attractive solution to this puzzle. The occurrence of ‘Micaëla knows’ in (1) demands less than does the same phrase in (2). The mention of the possibility that Moralès is lying ‘raises the standards’ for knowledge. An inference from (1) and $\sim(2)$, in their natural contexts, to (3) is therefore equivocal. (1) is true in

¹ The delightfully apt ‘abominable’ for conjunctions such as these is from DeRose (1995), pp. 27-28.

its natural context, and $\sim(2)$ is true in its, but in neither context are both true, so in neither context—nor, apparently, in any other²—is (3) true.

It is controversial whether contextualism about ‘knows’ provides a satisfactory solution to the puzzle, and indeed whether it is plausible for other reasons.³ But it is not difficult at least to feel the pull of ‘knows’ contextualism as an appealing response to this sort of data. For myself, I believe that the contextualist treatment of this puzzle is correct. But a sustained defense of that project is not my present concern. (See Ichikawa, forthcoming.)

What I would like to argue in this paper is that the considerations favoring contextualism about knowledge also favor contextualism about counterfactuals. The skeptical paradox given above has a parallel puzzle for counterfactuals, and the contextualist solution is attractive and plausible. Contextualism about counterfactuals, then, ought to be at least as appealing as contextualism about knowledge. Furthermore, parallel contextualist accounts of each can explain some apparent connections between knowledge and counterfactuals: for example, they can vindicate safety and sensitivity, construed as counterfactuals, as necessary conditions on knowledge.

I. The Counterfactual Puzzle

The Count insists that the closet be opened; the Countess reluctantly acquiesces, and out comes the maid. Both Count and Countess are surprised—the Count, because he jealously suspected his wife of hiding a man in the closet; the Countess, because she *did* hide a man in the closet. The Countess is privately glad that the man somehow swapped places with the maid. Why is she glad? Because, she truly says to herself:

4. If the Count had found a man, he would have become violent.

We may stipulate that the Count has the temperament that (4) seems to require: he is relevantly prone to violence. But even so, of course, it’s clear that it was *possible* for the Count to find a man without becoming violent. For instance, it could have been that the Count found a man, and the man immediately shot him dead. In that case, of course, the Count would not have become violent; he would not have had enough time. So (5) is false:

5. If the Count had found a man and the man had immediately shot him dead, the Count would have become violent.

² Of course, (3) may be true of some *situations*; suppose Micaëla sees that José is off duty and doesn’t know whether Moralès said he was. The point is that, holding the described situation fixed in the way stipulated, no conversational context delivers a true reading of (3).

³ Cohen (1998) argues that Lewis’s approach cannot solve the Gettier problem while also solving the lottery paradox; Stanley (2005) argues that contextualism of this sort wrongly predicts the acceptability of many standard abominable conjunctions. I respond on behalf of a Lewisian contextualism to these objections in my ‘Quantifiers and epistemic contextualism’ (forthcoming). Hawthorne (2004) argues that contextualists must deny attractive links between knowledge and action and assertion; similarly, Williamson (2005) argues that contextualists must embrace awkward links between knowledge and appropriate belief and assertion. These *prima facie* tensions between contextualism and ‘knowledge first’ views warrant much further attention, and comprise what I hope will be a fruitful future research project.

As before, the puzzle arises from juxtaposing our two verdicts.

6. If the Count had been shot dead by a man he found, he would not have become violent, but if he had found a man, he would have become violent.

Something is surely wrong; (6) is, under the circumstances, unacceptable, just as (3) was. What is to be done?⁴

One option is to reject the initially-attractive (4). It's not too hard to get into a frame of mind in which this is an attractive choice. One may reason: "if the Count had found a man, of course he *might* not have become violent, because he might have been shot dead instantly. So it can't be the case that he *would* have become violent." It's easy to see, however, that the grounds for rejecting (4) will generalize to almost all counterfactuals; almost always, although we might want to say that if A were the case, C would be the case, we can still think of a situation B, compossible with A, such that if A and B were the case, C would not be the case. If these are sufficient grounds for rejecting the initial counterfactual, than most ordinary counterfactuals would have to be judged false.⁵ This strategy, then, bears an obvious similarity to radical skepticism, on which the skeptical paradox is resolved by denying the initially-plausible knowledge attribution. This strategy, like radical skepticism, offers an extreme solution.

Another possible option is to admit that (5) is true. This strategy is unappealing to the point of absurdity. I mention it because its parallel position with regard to the skepticism puzzle—the

⁴ DeRose (1999) considers something like this puzzle; like me, he considers it in relation to the skeptical paradox. His version of the puzzle relies on might-counterfactuals. Adapting the examples to match mine, DeRose's version of the puzzle is:

- (i) If the Count had found a man, he would have become violent.
 - (ii) If the Count had found a man, he might not have become violent.
- So, (iii) If the count had found a man, he might not have become violent; nevertheless, he would have become violent if he had found a man.

DeRose observes that (i) and (ii) are very difficult to deny, but that (iii), which appears to follow, is very difficult to accept. DeRose's solution is to treat the might-counterfactual in (ii) and (iii) as expressing not the dual of the would-counterfactual, but instead the epistemic possibility of it. He then assimilates the unacceptability of (iii) to its Moore-paradoxical nature. (By contrast, he opts for a semantic treatment of the skeptical paradox.) Whatever the merits of this approach, it does not seem to resolve the version of the puzzle I have given, as none of (4)-(6) contains a might-counterfactual.

⁵ Alan Hájek (manuscript) suggests this approach. See also Hawthorne 2004, p. 5 fn. 10. Hájek's paper (and Hawthorne's comment) cites the apparent physical fact that indeterminacy governs much of the physical world; since (a prevailing interpretation of) quantum mechanics tells us that there is a genuine chance that some physical object spontaneously vaporizes at a given time, if I were to drop this cup, it might vaporize in midair; therefore, it is false that if I were to drop this cup, it would reach the ground. The role of quantum mechanics in this argument is to support the might-counterfactual, which is (assumed to be) inconsistent with the corresponding would-counterfactual. But, in a great many cases, the might-counterfactuals are independently plausible; quantum mechanics are inessential to the argument. In particular, there is no assumption, in the argument against (4) sketched here, that among the *nearest* worlds where the Count finds a man are worlds where the Count finds a man who instantly shoots him dead. And at any rate, as the previous footnote noted, my argument does not rely on might-counterfactuals as premises. Even if we suppose that the world is deterministic and that the worlds where a man in the closet shoots the Count dead are all more distant than the worlds where a man is in the closet, but does not shoot the Count dead, the argument from (6) and the negation of (5) to the negation of (4) seems cogent.

view that Micaëla does know that Moralès isn't lying, and in general, that we know skeptical scenarios to be false—is popular.⁶

A third solution corresponds to knowledge contextualism: we say that (4) expresses a truth in its natural context, and that (5) expresses a falsehood in its natural context, but that their natural contexts do not overlap, and so the inference from (4) and the negation of (5) to (6) is invalid. In its context, (6) is false.

What could a contextualist account of counterfactuals along these lines look like? Let us begin by considering, as a point of comparison, David Lewis's (somewhat) invariantist semantics for counterfactuals. (I shall explain shortly why it is not quite accurate to describe Lewis's view as entirely invariantist.)

II. David Lewis on Counterfactuals

According to David Lewis's influential (1973) account, counterfactuals are *variably strict conditionals*. A counterfactual is true just in case some material conditional holds in all of the relevant set of worlds. This is an appealing shape for a view of counterfactuals; my own view develops this observation in a contextualist way: which worlds are relevant depends in part on the conversational context of the speaker. But Lewis develops his variably strict conditional account in a different way. According to Lewis, the relevant set of worlds is the set of worlds where the antecedent is true that are *most similar* to the actual world—or, more precisely, more similar to the actual world than any worlds where the antecedent is false.⁷ So, to evaluate *if A were the case, then C would be the case*, we examine the sphere of A-worlds that differ less than any non-A-worlds from the actual world, and check whether C is true in all of those worlds.

This is meant to explain some apparent logical features of counterfactuals. For instance, it explains why counterfactuals are not closed under strengthening of the antecedent, the way that the material conditional is:

⁶ If, like me, you are convinced that the parallel nature of our two puzzles—the one about knowledge, and the one about counterfactuals—puts pressure in favor of parallel solutions, this will provide you with a reason to be uneasy about the widespread nonskeptical invariantist treatment of the skeptical paradox.

⁷ Lewis himself gives this disjunctive analysis: (D) the counterfactual is true just in case either there are no A-worlds, or some (A & C)-world is nearer than all (A & ~C)-worlds. My quantified formulation is (Q): all the nearest A-worlds are C-worlds. Given the limit assumption, my (Q) and Lewis's (D) are equivalent. (Lewis rejects the limit assumption, so he will deny the equivalence—I'm grateful here to Lee Walters. For present purposes, the limit assumption is a harmless simplifying assumption.) Here is a proof of the equivalence, given the limit assumption: Suppose (Q) is true. All the nearest A-worlds are C-worlds. If this is true vacuously, because there are no A-worlds, then (D) is true in virtue of its first disjunct. If (Q) is non-vacuously true, then there are some A-worlds, and all of the nearest of them are (A & C)-worlds. For any one of them, *w*, it is an (A & C) world that is closer than all (A & ~C) worlds—no (A & ~C) world can be closer than *w* because *w* is among the nearest worlds, and no (A & ~C) world can be tied with *w* for closeness because by stipulation *all* the nearest worlds are C worlds. So my version entails Lewis's.

Now suppose not-Q; that is, among the nearest A-worlds are some ~C-worlds. Since this entails that there is at least one (A & ~C)-world, the first disjunct of (D) is false. And since these (A & ~C)-worlds are among the nearest worlds, there are no worlds nearer than them. In particular, there are no (A & C)-worlds nearer than them; they are therefore counterexamples to the second disjunct of (D). So not-(Q) entails not-(D); Lewis's version entails mine.

$$A \supset C \Rightarrow (A \& B) \supset C$$

But not:

$$A \Box \rightarrow C \Rightarrow (A \& B) \Box \rightarrow C$$

My (4) and (5) provide an apparent counterexample to the latter: if the Count had found a man, he would have become violent, but if he'd found a man and been shot dead, he would not have become violent. Lewis's explanation is that, although both counterfactuals are strict conditionals, they are strict conditionals over different domains. The first counterfactual's domain includes the nearest worlds where the Count finds a man—these are worlds in which the count finds the mild-mannered and not-at-all-violent Cherubino, who was until recently hiding in the closet. The second counterfactual's domain involves a different, more distant set of worlds. Lewis comfortably handles these data.

Lewis runs into trouble, of course, with the puzzle given above; he gives the intuitively right verdict about (4) and (5), but offers no means to resist the unwholesome (6), the apparent conjunction of those intuitive verdicts.

A famous objection to Lewis's view focused on Lewis's invocation of similarity. Kit Fine (1975) observed that sometimes, counterfactual possibilities would have had dramatic effects on the course of history. So, for instance, *had Nixon pressed the button (to launch nuclear missiles), there would have been a nuclear war*. This counterfactual is very plausibly true—that's why we're glad he didn't press it—but it appears that Lewis's account will return the verdict that it is false. Consider the world where he presses the button, but the button malfunctions, and history runs its course as per actuality. This world is more similar to the actual world than the Armageddon world; so Lewis's account would consider this one the relevant world for evaluation of the counterfactual.

Lewis's (1986) response is that his invocation of 'similarity' is not to be thought of as overall macro-level similarity, but something more complicated—he gives a baroque system of weighted rules, outlining how violations of our actual laws, of various sizes, negatively affect similarity. As a preface to this project, however, Lewis admits that he is attempting to articulate only one notion of similarity that we make use of in evaluating counterfactuals; some counterfactuals call for different similarity relations.

Of particular interest for my project is the suggestion that among the factors that influence what kind of similarity relation is relevant are facts about the conversational context. (Lewis uses the unfortunate term 'vagueness' in a way which can here only be understood to mean 'context-sensitivity'.)

Lewis (1986) writes:

What is going on, I suggest, can best be explained as follows. (1) Counterfactuals are infected with vagueness, as everyone agrees. Different ways of (partly) resolving the vagueness are appropriate in different contexts. Remember the case of Caesar in Korea: had he been in command, would he have used the atom bomb? Or would he have used catapults? It is right to say either, though not to say both together. Each is true under a resolution of vagueness

appropriate to some contexts. (2) We ordinarily resolve the vagueness of counterfactuals in such a way that counterfactual dependence is asymmetric ... (3) Some special contexts favor a different resolution of vagueness, one under which the past depends counterfactually on the present and some back-tracking arguments are correct. If someone propounds a back-tracking argument, for instance, his cooperative partners in conversation will switch to a resolution that gives him a chance to be right. ... But when the need for a special resolution of vagueness comes to an end, the standard resolution returns. (4) A counterfactual saying that the past would be different if the present were somehow different may come out true under the special resolution of its vagueness, but false under the standard resolution. If so, call it a *back-tracking counterfactual*. Taken out of context, it will not be clearly true or clearly false. Although we tend to favor the standard resolution, we also charitably tend to favor a resolution which gives the sentence under consideration a chance of truth. (34)

It is easy to underestimate the significance of Lewis's concession here: Lewis is endorsing a limited form of contextualism about counterfactuals. Which notion of similarity is relevantly salient may change according to the conversational context; therefore, whether a counterfactual sentence is true depends not only on the antecedent, the consequent, and the state of the world, but also upon the conversational context. When faced with a counterfactual, we imagine a scenario described by the antecedent. But just what we are to imagine is underdetermined by the antecedent alone; *how* are we to imagine Caesar in command in Korea? Should we imagine him having replaced General MacArthur? Or should we imagine the Roman Empire having expanded to the east end of Asia?

The same question underwrites questions about backtracking counterfactuals. Shy Stan says, 'I'm terrified of public speaking. If I were the White House spokesman, I'd give the worst press conferences.' Quick Kevin retorts: 'You'd never be White House spokesman unless you got over your public speaking difficulties, Stan. If you were the White House spokesman, you'd give great press conferences (because you'd've had to have gotten over your stage fright).' Each may well be expressing a truth—even though the truths take the form of apparent contradictions.

Indicative conditionals may be understood to work in much the same way—the indicative mood signals an alternate way of filling out the scenario. (See McCawley, 1996, 90-91.) Indeed, in some circumstances, counterfactuals can even function in much the same way that indicative conditionals do. Here, I borrow from Dorothy Edgington (manuscript): You and I team up on a treasure hunt, and the organizer has given me a hint: 'it's either in the attic or the kitchen.' We split up; I search the attic, and you search the kitchen. I find the treasure in the attic, then you ask me why I sent you to the kitchen. I reply: 'Because if it hadn't been in the attic, it would've been in the kitchen.' I here speak truly, even if the organizer's second choice for a hiding place was the basement.

Once we are convinced that this sort of limited contextualism is true about counterfactuals, it is worth investigating whether such context-sensitivity runs much deeper—and what sorts of general principles might underwrite all counterfactuals.

III. Counterfactual Contextualism

Let's return to our puzzle about counterfactuals:

4. If the Count had found a man, he would have become violent.
5. If the Count had found a man and the man had immediately shot him dead, the count would not have become violent.
6. If the Count had been shot dead by a man he found, he would not have become violent, but if he had found a man, he would have become violent.

The knowledge contextualist resolves his puzzle by suggesting that ‘know’ refers to a different relation in each premise; the counterfactual contextualist, similarly, admits that (4) and (5) each express a truth, but that changing features of the conversational context prevent the apparent entailment from (4) (in its context) and $\sim(5)$ (in its context) to (6) (in its context).

One bit of support for this approach is that pairs of counterfactuals like these appear to be sensitive to dynamic conversational features. Pairs of counterfactuals like (4) and the negation of (5) *can* be acceptably conjoined in some conversational contexts, called Sobel sequences. Consider this somewhat more palatable way of putting something like (6):

7. If the count had found a man, he would have become violent. But, of course, if he’d found a man and he immediately shot him dead, he would not have become violent.

This bit of speech is comprehensible. Things change dramatically, however, for *Reverse* Sobel sequences. Suppose we put the counterfactuals into the opposite order:

8. If the count had found a man who immediately shot him dead, he would not have become violent. But, of course, if he’d found a man, he would have become violent.

Something is deeply wrong with (8). This conjunction has the appearance of a flat contradiction. This asymmetry has convinced some philosophers and linguists that there must be a dynamic element to counterfactuals. (See Gillies 2007; McCawley 1997.)

Incidentally, the same pattern emerges for knowledge attributions; (9) is preferable to (10):

9. Micaëla knows that José is off duty, but of course, she doesn’t know whether Moralès is lying.⁸
10. Micaëla doesn’t know whether Moralès is lying, but she does know that José is off duty.

My approach to knowledge is broadly Lewisian;⁹ S knows that p only if S’s evidence eliminates *all* (and this is a context-sensitive restricted quantifier) the not-p cases; the domain of that quantifier depends in part on the speaker’s context. So the asymmetry between (9) and (10) can be explained thus: to say Micaëla knows José is off-duty is to say that all the (relevant) José-on-

⁸ Sentence (9), and its parallel (7), might intuitively be thought to offer a sort of a contradiction: the first conjunct is asserted, then retracted with consideration of the second conjunct. Depending on the particulars of the contextualist view, one might respect that intuition, by treating the conjunction as false, with the second conjunct implicitly cancelling the first: “A, but B, so I guess not-A after all.” Alternatively, one might propose a form of contextualism on which both conjuncts are true, with the relevant contextual parameter shifting mid-sentence. Even on the former view, where (7) and (8), and (9) and (10), are both false, an asymmetry remains: (7) and (9) are much more comprehensible than are (8) and (10).

⁹ Lewis (1996).

duty possibilities are eliminated. Once we're taking seriously possibilities in which Moralès is lying, those possibilities become relevant, and the knowledge attribution is no longer apt.

Contra Lewis, the mere fact that one is attending to a possibility may not be enough to force it into relevance; it may be that one must be attending to a possibility in a certain kind of way—roughly, a *taking-seriously*.¹⁰ Furthermore, even once a possibility is included in the domain, it may be easier to retreat to nonskeptical contexts than Lewis supposed. One way to do this might be explicitly to rule out the possibility: 'If I were a brain in a vat, I wouldn't have hands. *But that's silly; of course I'm not a brain in a vat.* I know I have hands.' If this sort of move is legitimate here, it may also be in the case of the counterfactual: 'If the Count had found a man, he could have been instantly killed. *But of course that's silly; Cherubino wouldn't have killed the Count.* If the Count had found a man, he would have become violent.' (A minority of philosophers I speak to report that, contrary to my starting datum, utterances of sentences like (6) are unproblematic. It may be that these philosophers are charitably attributing such a 'but of course that's silly' cancellation.)

This parallel suggests an account of counterfactuals:

If A were the case, C would be the case is true just in case all of the A possibilities are C possibilities.

This 'all', like the one in Lewis's account of knowledge, takes a context-sensitive scope. So, we might say, following Lewis:

If A were the case, C would be the case is true just in case all of the A possibilities are C possibilities (Psst!—except those possibilities we're properly ignoring).

This account is an example of the type laid out by McCawley (1996, 91).

Although Lewis (1996) provided the groundwork for a strong contextualist solution to the knowledge puzzle, he did not seem to think much of the corresponding solution to the counterfactuals puzzle, opting instead (1986) to further complicate his static similarity relation in an attempt to capture our ordinary judgments (even while admitting that he was attempting to capture only the 'standard resolution' for counterfactuals). The Lewis of 1973 called a strict conditionals view of counterfactuals in which the relevant class of worlds is contextually determined, 'defeatist', and accused such an approach as 'consign[ing] to the wastebasket of contextually resolved vagueness something much more amenable to systematic analysis than most of the rest of the mess in that wastebasket.' (13) It's not clear whether the Lewis of 1996 retained his contempt for attribution of context-sensitivity; was Lewis really willing to throw *knowledge* into that 'wastebasket', when counterfactuals were too good for it?¹¹ Lewis's own

¹⁰ 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) for further critique of Lewis's version of the rule of attention.

¹¹ I have discovered in conversation that some philosophers are sometimes leery of accepting contextualism about counterfactuals because they think that counterfactuals have a robust role to play in explaining the world—perhaps, as Lewis thought, causation is understood in terms of counterfactuals; perhaps laws are. However, contextualism about counterfactuals does not by itself rule out a robust role for counterfactuals in scientific, philosophical, or other

reservations notwithstanding, a contextually variable strict conditionals analysis of counterfactuals is every bit as plausible as a contextually variable infallibilist analysis of knowledge.

It is such a view that I wish to develop.

IV. All Possibilities

Lewis's treatment of knowledge provides an instructive analogy. He gives a contextualist account of knowledge in terms of the context-sensitive 'all cases': S knows that p just in case S's evidence eliminates all the not- p cases. He supplements this view with a characterization of which cases count as part of 'all cases' in a given situation evaluated in a given conversational context. It is worthwhile to note, however, that the proposed treatment of 'knowledge' in terms of the context-sensitive 'all cases' is logically independent from Lewis's particular rule-based account of the latter. In my own view, Lewis's proposed rules are probably no better than approximately correct; this should not distract us from the well-motivated approach to knowledge.¹²

My suggestion is that the context-sensitive restricted 'all cases' invoked in Lewis's account of knowledge is the same one that is invoked in my treatment of counterfactuals above. Knowledge attributions and counterfactual statements are different generalizations over the same sets of possibilities: when I say I know that p , I say that my evidence eliminates all the not- p cases in the set. When I say that if A were the case, C would be the case, I say that the material conditional, $A \supset C$, holds in all the cases in the set.

One rule that Lewis does seem to have gotten right is the 'Rule of Actuality,' which ensured knowledge its factivity; the world of the subject was always included in the relevant domain, so the subject could never know that p without its being the case that p . The same rule, applied to my approach to counterfactuals, guarantees that counterfactual *modus ponens* be truth-preserving. If the actual world is always among the relevant worlds, then the counterfactual demands that $A \supset C$ be true; if A is true, then, so is C .

It is plausible that something in the ballpark of Lewis's Rule of Attention is operative in the metasemantic story about how conversational context fixes the relevant domains; this is how Lewis resolves the skeptical paradox. When we utter (1), we properly ignore the distant possibility in which Moralès is lying to Micaëla about José; the knowledge attribution comes out true: Micaëla knows that José is off-duty, because her evidence (the testimony) rules out *all* of the cases where José is on-duty, where that *all* is a restricted quantifier that does not include in its domain cases in which Moralès is lying. But, when we consider sentence (2), we do not ignore

inquiry. On my view, counterfactuals, evaluated at a context, do express objective facts about the world; this is in no tension with contextualism; it is also an objective fact that I am here now, even though 'I,' 'here,' and 'now' are context-sensitive. (But for my own part, I am skeptical about counterfactual analyses of just about anything; counterfactuals seem too sensitive to irrelevant features of the world—finking and masking problems for counterfactual analyses of dispositions provide an obvious example of the problem; Ichikawa (2009) offers a similar objection to Timothy Williamson's (2007) counterfactual treatment of philosophical intuitions.)

¹² These ideas are developed in my (forthcoming).

the possibility that Moralès is lying; once that possibility is sufficiently salient, the relevant domain expands to include it. Since Micaëla's evidence doesn't eliminate this possibility, she doesn't know that it does not obtain. And, relative to this new context, in which the lying-Moralès possibility is included as relevant, (1) expresses a falsehood; this is why (3) is not warranted. The principal difference between the (original) context of (1) and the context of (2) is that we were ignoring certain possibilities when we considered (1); this rendered them irrelevant with respect to the evaluation of the universal, and likewise the knowledge attribution. (As indicated above, one needn't, in going this far, agree with Lewis that mere mention of a possibility *forces* it into the domain, and can only be removed by the passage of time spent ignoring it.)

With something (admittedly imprecisely) like a Rule of Attention for 'all cases', and the account of counterfactuals in terms of 'all cases', the counterfactual puzzle is resolved in the same way as the skeptical puzzle: when we say (4), we say that all (and here we restrict the quantifier, ruling out irrelevant cases, such as the very unusual cases where the man shoots the Count dead, as properly to be ignored) cases in which the Count finds a man are cases where the Count becomes violent. But once the possibility of the Count finding a man and immediately being killed is raised, it is no longer properly ignored. This because, at least, it is no longer ignored at all. Since this is a possibility in which the Count does find a man but doesn't become violent, (5) is false as uttered. Note also that this dynamic feature of counterfactual discourse explains why (4) is unacceptable once (5) has been explicitly denied, even though it was acceptable on its own. Once the possibility that the Count finds a man and is shot dead is put on the table, both universals, (4) and (5), are false. (The same pattern holds for (1) and (2); in a context in which the possibility of misleading testimony is being taken seriously, (1) is false.)

V. Relating Knowledge and Counterfactuals

I have suggested, but not argued, that the same restricted domains apply in both the knowledge case and the counterfactual case. Of course, this needn't be so; one might have contextualist views of knowledge and counterfactuals without the proposed identity. Why think that knowledge attributions and counterfactual utterances express universal quantifiers over the very same domains? The strongest case in favor of this identity comes in the form of particular cases: consideration of cases that are counterexamples to a knowledge generalization, when those cases are also counterexamples to the counterfactual generalization, undermine the knowledge attribution and the counterfactual alike. Blanche is, and has long been, Professor of Abstract Science. This, as Melissa is well aware, won't change any time soon. So (11) is true:

11. Melissa knows that Blanche will be Professor of Abstract Science tomorrow.

In the unlikely event that Ida, the principal, abandoned her position tonight, Blanche would leave her post as Professor of Abstract Science and take her place:

12. If Ida resigned tonight, Blanche would be principal tomorrow.

The domain for (11) and (12), in their natural, 'nonskeptical' context, does not include cases in which Blanche resigns tonight. (These cases are very distant and legitimately ignored; Blanche is not likely to leave the university under any circumstances.) But of course, she *could* resign

tonight. We may explicitly bring those cases into the domain; doing so undermines both sentences:

13. Of course, Blanche could resign tonight, but Melissa knows that Blanche will be Professor of Abstract Science tomorrow.
14. Of course, Blanche could resign tonight, but if Ida resigned tonight, Blanche would be principal tomorrow.

Both of these sentences are false in their contexts—each is undermined in just the same way, by the consideration of a particular counterexample: a case where Blanche resigns tonight. It is a counterexample to the knowledge attribution because it's a case where Melissa's evidence is as it actually is, but Blanche will not be Professor of Abstract Science tomorrow; it is a counterexample to the counterfactual utterance because it's a case where Ida resigns tonight, but Blanche is not principal tomorrow.

I am unaware of any data suggesting that the two domains should come apart; consideration of particular cases like this one provides some reason for thinking that they will shift together.¹³ General theoretical simplicity provides a second reason. A third positive reason to treat the relevant domains as coextensive is that doing so permits attractive explanations for the apparent necessary conditions of safety and sensitivity on knowledge.

I turn now to these explanations, beginning with sensitivity.

¹³ A referee suggests the following disanalogy: knowledge attributions seem sensitive to *stakes* in a way that counterfactuals are not. The more important the propositional content of the knowledge attribution is, the more it takes to count as 'knowing'. See DeRose (1992) and Stanley (2005) for articulations of stakes-sensitivity for knowledge-attributions, or for knowledge, respectively. (See also Cohen (1999) for the former.) I am not convinced that this represents a disanalogy that undermines my claims about counterfactuals. First, consider that it is not obviously the stakes *per se* that are playing a relevant shifting role in DeRose, Cohen, and Stanley's airport and bank cases; it may be instead that the high stakes of the situation cause us to attend to more skeptical possibilities: since it is very important whether the bank is open, I dwell on more ways in which I might go wrong, thus entering those cases into the domain in something like the standard Lewisian way. (Compare: if I have a very important presentation to give, I may think about more ways in which things might go wrong: "What if I mispronounce this name? What if nobody shows up? What if I sweat through my shirt?") If this is the way that high stakes expand domains for knowledge attributions, then it should not be too surprising if stakes turn out to play a diminished role in the evaluation of counterfactuals, since truths about counterfactual situations are much less likely to be of great practical importance in deliberation. It is very important whether the bank will be open tomorrow; it is of relatively trivial importance, for the purpose of what to do with my check today, whether the bank would be open tomorrow, were some counterfactual supposition to obtain. So perhaps we shouldn't expect the effect of stakes on counterfactuals to be as widespread as in the case of knowledge. In some cases, however, it is plausible that high stakes result in expanding the domains for knowledge and counterfactuals alike. Suppose we know that Morgan decided the bank hours last year, and that he decided it should be open until noon on Saturdays. In a relatively lax, 'nonskeptical' context, we speak truly on Friday in saying both "I know the bank will be open until noon tomorrow" and "if Morgan had decided to close the bank on Saturdays, the bank would be closed tomorrow." Now suppose it's of critical importance whether the bank will be open tomorrow. The high stakes lead us to think about the possibility that the bank has changed its hours since last year, which is why the knowledge attribution is unacceptable. In such a case, the counterfactual may also be unacceptable, since we're treating as live the possibility that the bank has recently changed its hours. The knowledge attribution and the counterfactual are undermined in the same way. (One way we might think of things: we're undermining the regularity between Morgan's decision and the hours this week. That is, we're breaking the connection between the evidence and the knowledge, in one case, and between the antecedent and the consequent, in another.)

VI. Knowledge and Sensitivity

Although particular attempts to develop the idea have proven susceptible to apparent counterexample, many epistemologists have been attracted to the idea that a certain counterfactual relation between a subject and a proposition is a necessary condition for knowledge, by that subject, of that proposition. Robert Nozick (1981), for instance, proposed a *sensitivity* requirement on knowledge: S knows that *p* only if, were *p* not the case, S would not believe that *p*. As many critics have observed, this requirement appears problematic—for instance, combined with anti-skepticism, it appears to entail the failure of single-premise closure. (I know that I see Laura—if I didn't, it wouldn't appear to me that I did; but I don't know that I don't see a perfect imposter duplicate Laura instead of the genuine article—if I did, I'd still think that I saw Laura.) Nevertheless, the pre-theoretic appeal of the sensitivity condition on knowledge must be explained. There is something cogent about the argument: *how can you know that p? You'd think that p even if p weren't the case!*

Many of the influential objections to sensitivity as a necessary condition on knowledge have assumed a standard (invariantist) Lewisian semantics for counterfactuals. On my preferred contextualist account of counterfactuals, along with (a suitable development of) Lewis's contextualist account of knowledge, sensitivity, understood in its standard counterfactual formulation, is a genuinely necessary condition on knowledge.¹⁴ (In §8 below, I defend the claim that the counterfactual formulation of safety is also a necessary condition on knowledge.)

Here is my preferred account of knowledge:

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.¹⁵

(Understand this statement equivalence as true in all contexts.) I clarify and defend this view in my (forthcoming). What is important for present purposes is that it is contextualist in Lewis's spirit: it develops Lewis's account by adding a basing requirement; conjunct (ii) in my formulation is Lewis's picture of knowledge, and the 'all' contained in it is the contextualist restricted quantifier Lewis uses in his approach to knowledge, and that I additionally use in my approach to counterfactuals.

Let us gloss sensitivity in the following way:

A belief that *p* on the basis of E is sensitive just in case, if *p* hadn't been the case, the subject wouldn't have believed that *p* of the basis of E.

This is my gloss of sensitivity, anyway; sensitivity is sometimes presented without the invocation of the subject's basis for belief: S's belief that *p* is sensitive just in case, were *p* not the case, S wouldn't believe that *p*. I invoke the basis with an eye toward cases in which a subject believes

¹⁴ This is not to commit to any sort of JTBX-style analysis of knowledge; I do not think that the correct theory of knowledge will include sensitivity as a conjunct. The claim is only the literal one: knowledge entails sensitivity; necessarily, all cases of knowledge are cases of sensitive belief. See Williamson (2002, 2-5).

¹⁵ This is meant as a general account, covering both basic and nonbasic knowledge. Basic knowledge is knowledge that is believed on the basis of itself.

that p on good grounds, but would believe that p on bad grounds, were p not the case.¹⁶ Here are two examples:

- (a) A wife believes, on the basis of excellent evidence, the truth that her husband is faithful. But she is psychologically unable to face difficult truths. So, had her husband been unfaithful, she would have self-deceived herself into irrationally believing him to be faithful.
- (b) The priest has gone to see whether the god is on the mountain; he climbs the mountain, sees the god, and comes to believe that the god is on the mountain. The god, anxious for his subjects to believe him to be on the mountain, plants a trap whenever he goes away, which will zap any visitors' brains. Their memories of having been to the mountain are erased, and they are given an unreflective belief that the god is on the mountain. So, had the god not been on the mountain, the priest would have had his memory wiped and been left with the unreflective belief that the god is on the mountain.

The properly-formed beliefs of the wife and the priest, intuitively knowledge, are sensitive in my sense.

A sensitivity requirement occupies a central role in Dretske's and Nozick's treatments of knowledge. But sensitivity is often thought to be problematic as a requirement for knowledge. As I pointed out above, sensitivity, combined with anti-skepticism, appears to imply the failure of single-premise closure. And Ernest Sosa, among others, has provided, independent of this worry, apparent counterexamples to a sensitivity requirement for knowledge. Here is an influential one:

On my way to the elevator I release a trash bag down the chute from my high rise condo. Presumably I know my bag will soon be in the basement. But what if, having been released, it still (incredibly) were not to arrive there? That presumably would be because it had been snagged somehow in the chute on the way down (an incredibly rare occurrence), or some such happenstance. But none such could affect my predictive belief as I release it, so I would still predict that the bag would soon arrive in the basement. My belief seems not to be sensitive, therefore, but constitutes knowledge anyhow, and can correctly be said to do so. (Sosa, 1999, 145-6)

The garbage chute case appears to be a case of insensitive knowledge. In this case, the closest worlds in which the bag won't make it to the basement are worlds in which the subject falsely believes, on the basis of his actual evidence, that the bag will make it to the basement. So, Sosa has provided a counterexample to this principle:

Target: S knows that p on the basis of E only if, in the nearest possible worlds in which not- p , S does not believe that p on the basis of E.

Is the garbage chute case a counterexample to sensitivity as a requirement for knowledge? Recall that sensitivity invokes a counterfactual: S's belief that p on the basis of E is sensitive just in case, if p hadn't been the case, S wouldn't have believed that p on the basis of E.

¹⁶ Nozick himself (1981) makes a move similar to mine, relativising sensitivity to methods instead of bases.

A sensitivity requirement for knowledge is equivalent to Target if something sufficiently similar to Lewis’s semantics for counterfactuals is correct. Since Lewis’s semantics is widely assumed to be correct in much of the relevant literature, it is not surprising that a counterexample to Target should often be assumed without argument to be a counterexample to a sensitivity requirement.¹⁷ But an alternative semantics—one we’ve seen to be in some ways preferable—is now on the table. How does sensitivity fare if we understand counterfactuals, and knowledge, as described above?

I have placed the bag into the chute, and now know, on the basis of my experience, that it will shortly arrive in the basement. That is to say, my evidence rules out all the possibilities where the bag won’t shortly arrive in the basement (excepting, as always, the ones we’re properly ignoring). What of the counterfactual that would establish sensitivity? *If the bag were not to arrive shortly in the basement, I would not believe that the bag would arrive shortly in the basement.* This is true just in case all the possibilities in which the bag won’t land in the basement (again, with a suitably restricted ‘all’) are cases in which I don’t believe the bag will land in the basement. This is false; the relevant counterexamples are the ones Sosa discusses—cases where the bag, unbeknownst to me, gets stuck on the way down. So doesn’t Sosa’s counterexample to sensitivity generalize even to my version of the sensitivity counterfactual?

Perhaps not. On my approach, there is a context-sensitive element to both knowledge attributions and counterfactuals—the evaluation of each depends on which possibilities we’re properly ignoring. And the cases that falsify the counterfactual—the ones where the bag gets stuck in the chute—are exactly the ones we must properly ignore in order to go through with the knowledge attribution. The two claims that together appear to violate sensitivity, then, only hold if we switch contexts between them. When we consider the counterfactual, *what if the bag were to fail to make it to the basement?*, we expand the class of relevant possibilities to include some in which the antecedent holds. Note that, once we’ve done this, we can no longer make the knowledge attribution; it is defective to say, ‘if the bag weren’t to make it to the basement, I’d falsely believe the bag would make it to the basement. But I know the bag will make it to the basement.’

This is a common pattern with context-sensitive discourse; here is an instructive parallel. Consider this Extremely Attractive Principle:

EAP. All Φ s are F only if all G Φ s are F.

Someone sets out to provide a counterexample to the extremely attractive principle. He points to the Hermès store, filled with designer handbags. The cheapest handbag there costs more than \$2,000. ‘Everything in this store is expensive,’ he says, speaking truly. Then he adds, thinking of the wastebaskets, which contain such things as paper receipts: ‘The *garbage* in this store is not expensive.’ He speaks truly again. Then, conjoining his two true utterances, he claims to have found a counterexample to EAP—everything in the store is expensive, even though the garbage in the store is not expensive.

¹⁷ But see fn. 22 below for why Sosa cannot adopt all features of Lewis’s account.

The fallacy is transparent. The domain of the quantifier ‘everything’ varies according to the conversational context. We may truly say that everything in the store is expensive only when we are restricting our quantifier to the merchandise—only when we are properly ignoring things like the contents of the wastebaskets; when we say that the garbage in the store is not expensive, we stop ignoring the garbage, so the domain broadens to include it. This is why we cannot give the two sentences in the opposite order: *The garbage in the store is not expensive, but everything in the store is expensive*. The imagined counterexample relies on illicit changes in conversational context. This is clearly the right response in this case—we do not give up the claim that everything in the store is expensive only if there are no groups of things that are inexpensive. That ‘everything in the store is expensive’ plausibly expresses a truth only when wastebasket contents are ignored; considering whether the wastebasket contents in the store are expensive demands us to consider the wastebasket contents—once we do, we admit that they’re not expensive (and are no longer in a position to say “everything in the store is expensive”).

In just the same way, the garbage chute case needn’t threaten a sensitivity requirement for knowledge; that the subject knows the garbage will be in the basement is only plausible when the garbage-getting-stuck cases are ignored; considering whether the relevant counterfactual (*if the garbage weren’t to reach the basement, the subject wouldn’t believe that it would*) is true demands us to consider the garbage-getting-stuck cases—once we do, we admit that the counterfactual is false (and are no longer in a position to utter ‘the subject knows that the garbage will reach the basement’).

Although this sort of case does not threaten the truth of sensitivity as a necessary condition on knowledge, it does have the implication that the consideration of sensitivity counterfactuals is not in general an effective way to *evaluate* the truth of knowledge claims, since the context so easily shifts between the knowledge attribution and the corresponding sensitivity counterfactual. That the counterfactual, considered in its context, is false, does not necessarily tell against the knowledge attribution, considered against a more relaxed set of possibilities.¹⁸

I have so far defended the sensitivity requirement against apparent counterexamples by pointing out that, when counterfactual utterances like ‘if the bag weren’t to make it to the basement, you’d believe it to be in the basement’ are made, the context shifts to a skeptical one where the counterfactual is true, and where a knowledge claim would be false. I’ve defended the claim, therefore, that only in contexts in which one can truly assert the sensitivity counterfactual can one truly make the knowledge attribution. This is not yet to defend sensitivity as a necessary condition for knowledge. For this to be true, the sensitivity condition must apply in the contexts in which the knowledge attribution is true; it is not enough to show that the requirement holds in the cases in which the sensitivity counterfactual can be asserted (and in which the knowledge attribution is false). What of the nonskeptical contexts?

Consider the nonskeptical case. I’ve dropped the bag down the chute. My evidence eliminates all the cases where the bag doesn’t end up in the basement, so I know the bag is in the basement. (We’re ignoring cases where the bag gets stuck on the way down.) The sensitivity counterfactual, therefore, must be true, if sensitivity is a genuine necessary condition on knowledge. *If the bag*

¹⁸ Thanks to Jessica Brown here.

weren't to make it to the basement, I wouldn't believe, on the basis of E, that it would. This counterfactual sounds false, but I'm committed to its being true, in this nonskeptical context. Is this a decisive objection to my view?

No. Admittedly, the counterfactual sounds false, but my view explains and even predicts this: if it were uttered, it would expand the context to a skeptical one where it would be false. So of course it should *sound* false. Any time you heard it, it would be false. On my view, the counterfactual is true, even though the sentence would express a different, false counterfactual if uttered. It is true because all the antecedent-worlds are consequent-worlds: in all the worlds (in the domain) where the bag doesn't make it to the basement, it is not the case that I believe, on the basis of my actual evidence (*i.e.*, I put the bag in the chute, everything appeared normal, etc.), that the bag does make it to the basement. This is for the relatively trivial reason that there are no worlds in the domain where the bag doesn't make it to the basement and I have my actual evidence.

It is true that a certain unease remains about this counterfactual, even given my attempt to explain our reluctance away. After all, some propositions are easily recognized as true, even though they would be false if uttered. *I am not speaking* is often an example. The way to ameliorate this unease is, as before, to compare the situation with that that arises from the less controversial contextualist account of quantifiers.

All eighteen members of the department are in the seminar room. The chair looks around to ascertain whether they're ready to begin the meeting. 'Is someone in the library?' the chair wonders aloud. She counts the faculty present and concludes, 'no, no one is in the library; everyone is here.' She speaks truly. The domain includes the members of the department; *everybody is in the seminar room* and *no one is in the library*, because there are no department members at the library; they're all in the seminar room, right where they belong.

There is a librarian in the library, but he doesn't falsify the chair's assertion. He's outside the domain, properly ignored. An utterance of 'there is a librarian in the library' would, annoyingly, expand the domain to include him, and introduce a context in which 'everyone is here' is false and 'no one is in the library' is true.

The domain in the chair's conversation includes all department members; it does not include librarians. Quantifiers in this context ignore anything that isn't a department member. Therefore, it is true, in that context, that *all the librarians are in the seminar room* and that *no librarians are in the library*. There are no librarians in the library because nothing in the domain is both a librarian and in the library—because there are no librarians in the domain. (Every single thing in the domain—*i.e.*, every single department member—is a non-librarian.) And all the librarians are in the seminar room for just the same reason—there are no librarians.

These implications sound absurd, but they follow from well-supported contextualist theorizing about quantifiers. When we consider whether 'anything' is an X, when we know full well that X's are lurking around outside the domain, we expand the domain to include some X's and answer in the affirmative. We only think about the question in the more-inclusive domain. But it should be clear from the case above that in truth, relative to the initial, more restricted domain, it is true that no librarians are in the library, and that all the librarians are in the seminar room.

The sensitivity counterexample above had just the same structure: it is true, relative to the nonskeptical domain, that no cases in which the bag fails to reach the ground are cases in which the subjects believes it does on the basis of E. The cases that would falsify the universal are safely tucked away outside the domain—this even though at first mention of them, we shift our domain to include them. We should therefore not be too troubled by the counterintuitiveness of this result.

Sensitivity does not imply closure failure for knowledge. Apparent counterexamples equivocate. Here is the alleged consequence of sensitivity I began with:

I know that I see Laura—if I didn't, it wouldn't appear to me that I did; but I don't know that I don't see a perfect imposter duplicate Laura—if I did, I'd still think that I saw Laura.

The solution is straightforward and familiar: once we are taking seriously the imposter-Laura possibility, it is no longer true to say, 'I know that I see Laura'—and neither is it true to say 'if I didn't see Laura, I wouldn't think that I did.' So there is no counterexample to closure; deductive inference preserves knowledge. 'Knowledge' in the previous sentence must be understood univocally; it is true that sometimes we truly say sentences of the form 'S knows that p ', and also truly say 'S does not know that q ', even when S has correctly deduced q from p . These are cases in which the relevant contextual parameter is shifting between sentences.

VII. Strengthening the Antecedent

Let me pause here, before considering the safety counterfactual as a potential necessary condition on knowledge, to return to the apparent failure of counterfactuals to support strengthening the antecedent. It had been thought that (4) and not-(5) provided a counterexample; they are of the form $(A \Box \rightarrow C)$ and $\sim((A\&B) \Box \rightarrow C)$, respectively. However, given the contextualist machinery that has emerged, that there are sentences of those forms that each express truths in their natural contexts does not establish the failure of strengthening the antecedent. For strengthening the antecedent to hold, it must be the case that, *within a given context*, any time $(A \Box \rightarrow C)$ is true, so is $((A\&B) \Box \rightarrow C)$; since, as I have already emphasized, (4) expresses a truth only in contexts in which the possibilities discussed in (5) are ignored, this pair fails to provide a counterexample.

In fact, the validity of strengthening of the antecedent follows directly from the proposed view of counterfactuals in terms of quantifiers; $(A \Box \rightarrow C)$ is true just in case all the A cases are C cases; $((A\&B) \Box \rightarrow C)$ is true just in case all the A cases that are also B cases are C cases. In other words, the validity of strengthening the antecedent is just an instance of the 'extremely attractive principle' discussed in the previous section; apparent counterexamples equivocate in the same way that the bad argument about Hermès handbags and garbage did.^{19,20}

¹⁹ See Brogaard & Salerno (2008) for a defense of counterfactual strengthening the antecedent on similar (contextualist) grounds.

²⁰ An anonymous reviewer suggests that I am here relying too much on very fine context shifts: I claim that (4) and not-(5) fail to provide a counterexample to the validity of strengthening the antecedent, because the context in which (4) is true is a different one from that in which (5) is false. But this is only true if contexts are individuated very finely; shouldn't we allow a contextual parameter to shift between (4) and (5) within a single context?

VIII. Sensitivity and Safety

Sosa's argument, purportedly refuting a sensitivity requirement for knowledge, in fact refutes only the combination of a sensitivity requirement for knowledge and Lewis's semantics for counterfactuals.²¹ Since I have argued independently against the latter, Sosa's argument is not persuasive against a sensitivity requirement. This isn't yet to say that sensitivity *is* a requirement for knowledge—only that, for all Sosa has given us, there's no reason to take that possibility off the table. We may wonder whether to prefer a sensitivity requirement or an alternative, such as Sosa's *safety* requirement on knowledge. Here is Sosa's gloss on safety:

A belief is sensitive iff had it been false, S would not have held it, whereas a belief is *safe* iff S would not have held it without it being true. For short: S's belief $B(p)$ is sensitive iff $\sim p \Box \rightarrow \sim B(p)$, whereas S's belief is safe iff $B(p) \Box \rightarrow p$.²² (146)

As before, we should add a basing requirement to the gloss on safety. The modification is to protect safety from the same sorts of cases mentioned above—cases in which the subject responds properly to the evidence, but who could have behaved irrationally instead. Add to the two examples that the husband could easily have been unfaithful, or that the god could easily have been away, and we have failures of safety if we don't include a basis in the formulation, and cases of safety if we do. So we have:

S's belief that p on the basis of E is safe just in case, were S to believe that p on the basis of E , p would be the case.

In the garbage chute case, the relevant belief is safe because, although there are possible scenarios in which the subject believes as he does and does so falsely, none of those scenarios could *easily* obtain. The relevant worlds are too distant to undermine the safety counterfactual. Sosa thinks this to be a difference between safety and sensitivity, but I have argued that, in the cases where it is plausible that the subject knows that the garbage will make it to the basement, that belief is sensitive after all (although I admitted that, if we dwell on its sensitivity, the context

Suppose we do. The alleged counterexample's defect remains: the scope of the quantifiers in (4) and (5) are not the same. Perhaps we may say that the utterances exist in the same 'context' in some broader sense, but this does not save the argument from equivocation, since the contextual parameters vary between (4) and (5).

The reviewer suggests that it is sufficient for the invalidity of strengthening the antecedent that (4) be true and (5) be false. I do not think that this is the best way to think about validity. Compare the equivocal argument typed by this utterance: "he's a contextualist [pointing at Keith], therefore he's a contextualist [pointing at Tim]." The argument equivocates in the same way that the one above does (again, regardless of how one types contexts). And that the antecedent is true, while the consequent is false, doesn't impugn the validity of repetition. In the same way, that (4) is true, while (5) is false, doesn't impugn the validity of strengthening the antecedent. This is the same moral of that of the discussion in §6 of EAP, which is also valid, even though it is true that everything in the store is expensive, and false that the garbage in the store is expensive.

²¹ Or something sufficiently similar to it. See footnote 22 below.

²² Here, it is clear that Sosa is presupposing that $(A \ \& \ C)$ does not entail $(A \ \Box \rightarrow \ C)$. Otherwise, all true beliefs would be safe. Since this entailment is a theorem on Lewis's semantics, Sosa needs an account of counterfactuals that is both invariantist and not quite Lewisian. Perhaps Sosa could adopt an anticontextualist Lewisian framework for counterfactuals on which counterfactuals are strict quantifiers governing the closest antecedent worlds, while denying Lewis's axiom that the actual world is always uniquely closest to itself; if Sosa lets a sphere of worlds—the 'close' worlds—count as tied for closest, then he may be able to use Lewis's semantics for counterfactuals to obtain the results he needs.

will shift to one in which we must deny both the knowledge attribution and the claim to sensitivity).

In fact, we may draw a quite general relationship between sensitivity and safety. Sosa explains how he thinks they must differ: safety is the contraposition of sensitivity, and subjunctive conditionals, Sosa says, do not contrapose:

If water now flowed from your kitchen faucet, it would *not* then be the case that water so flowed while your main valve was closed. But the contrapositive of this true conditional is clearly false.
(152)

But, as before, my treatment of counterfactuals provides a debunking explanation of the apparent counterexample; when we consider the counterfactual Sosa gives, we consider a set of possibilities that doesn't include any cases where the water leaks through a closed valve; when we consider the contrapositive, *if water flowed from your faucet while the main valve was closed, water would not have flowed*, we turn to a new set of possibilities. So the case is consistent with the principle that a subjunctive conditional is equivalent to its contraposition, when we hold the context fixed—it's just that sometimes, one formulation forces us into a context in which the other no longer expresses a truth.

Don't forget the helpful analogy with context-sensitive quantifiers: (not everyone has studied semantics but) everyone understands English; of course, some Chinese people don't understand English. This does not refute the duality of the universal and the existential. Or, to tighten the cases, go back to the water faucet case and work with the quantifiers I use to gloss the counterfactuals:

15. If water were to leak, it would not leak from a closed faucet.
16. None of the possibilities where water leaks involve water leaking from a closed faucet.
17. If water were to leak from a closed faucet, water would not leak.
18. None of the possibilities where water leaks from a closed faucet involve water leaking.

Each of (16) and the negation of (18) can express a truth—(16), like (15), might be a useful thing to say if, for instance, one wanted to take steps to ensure that water wouldn't leak. Just make sure to close the faucet, and the water won't leak. We are ignoring cases in which the water manages to leak out of a closed faucet—our faucet is in good working condition, so these are reasonable cases to ignore. We stop ignoring them, though, when we mention them in (18) or its negation. The context changes, and our quantifiers take new scopes. (16) expresses a truth when uttered before the negation of (18), which also expresses a truth, but this is no threat to the duality of the universal and the existential, even though (16) and $\sim(18)$ appear to have the form of $(\forall x)\sim Fx$ and $(\exists x)Fx$.²³

On my treatment of counterfactuals, safety and sensitivity will be equivalent, when we hold fixed the context set. This is easily seen: a belief that p is sensitive just in case $\sim p \square \rightarrow \sim B(p, E)$. That is, just in case all the not- p possibilities are not- $B(p, E)$ possibilities; a belief that p is safe just in case $B(p, E) \square \rightarrow p$. That is, just in case all the $B(p, E)$ possibilities are p possibilities.

²³ This is a further example in favor of the point made in footnote 20.

Counterfactuals *do* contrapose, because universals contrapose, and counterfactuals are kinds of universals.²⁴

Safety and sensitivity are equivalent; furthermore, they are both genuine necessary conditions on knowledge. A belief that p , based on E , is sensitive just in case if p weren't the case, S wouldn't believe that p on the basis of E . That is to say, given my treatment of counterfactuals, a belief is sensitive just in case there are no possibilities in which S believes that p on the basis of E , but p is false.

Suppose S 's belief that p , based on E , is insensitive. Then there are some possibilities in which S believes p on the basis of E , where p is false. Since these are possibilities in which S believes p on the basis of E , these are possibilities where E obtains. So it follows directly from insensitivity that there are possibilities in which E and not p . On the proposed account of knowledge, S 's belief that p , based on E , is knowledge only if there are no possibilities in which E and not p . So, on the proposed accounts of knowledge and counterfactuals, knowledge entails sensitivity, which is equivalent to safety.

Conclusion

Knowledge and counterfactuals are intimately related; it is no accident that so many attempts have been made to understand the one in terms of the other. My suggestion is to understand both in terms of context-sensitive restricted universal quantifiers.

The motivation for a contextualist approach to knowledge generalizes to a contextualist approach to counterfactuals; each inherits its context-sensitivity from the context-sensitive 'all possibilities'. Understanding them this way explains the parallel nature of the skeptical paradox and the similar challenge against ordinary counterfactuals.²⁵

References

- Brogaard, B. and Salerno, J. (2008). Counterfactuals and context. *Analysis*, 68(1): 39-46.
- Cohen, S. (1998). Contextualist solutions to epistemological problems: Scepticism, Gettier, and the lottery. *Australasian Journal of Philosophy* 76(2): 289-306.
- Cohen, S. (1999). Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives* 13: *Epistemology*: 57-89.

²⁴ Brogaard and Salerno (2008) argue for counterfactual contraposition on contextualist grounds.

²⁵ A version of this paper was given at a colloquium at Indiana University in 2009; I am grateful to the participants there for a fruitful discussion. I am grateful also, for valuable conversations about counterfactuals and knowledge, and for comments on previous drafts of this project, to Jessica Brown, Dylan Dodd, Dorothy Edgington, Tamar Szabó Gendler, Alan Hájek, Doug Husack, Alex Jackson, Benjamin Jarvis, Karen Lewis, Barry Loewer, Sara Moss, Joshua Schechter, Martin Smith, Ernest Sosa, Jason Stanley, Brian Weatherson, Jonathan Weinberg, Crispin Wright, Elia Zardini, the Arché Basic Knowledge seminar, and an anonymous reviewer for *Philosophy and Phenomenological Research*.

- DeRose, K. (1992). Contextualism and knowledge attributions. *Philosophy and Phenomenological Research*, 52(4): 913-929.
- DeRose, K. (1995). Solving the skeptical problem. *The Philosophical Review*, 104(1): 1-52.
- DeRose, K. (2002). Assertion, knowledge and context. *The Philosophical Review*, 111(2): 167-203.
- Edgington, D. (manuscript). Do counterfactuals have truth conditions? Presented at the 2006 Eastern APA.
- Fine, K. (1975). Review of David Lewis's *Counterfactuals*, *Mind*, 84: 451-58.
- Gillies, A. (2007). Counterfactual scorekeeping. *Linguistics and Philosophy*, 30(3): 329-60.
- Hájek, A. (manuscript). Most counterfactuals are false. Unpublished manuscript.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. New York: Oxford University Press.
- Ichikawa, J. (2009). Knowing the intuition and knowing the counterfactual. *Philosophical Studies*, 145(3): 435-43.
- Ichikawa, J. (forthcoming). Quantifiers and epistemic contextualism. Forthcoming in *Philosophical Studies*.
- Lewis, D. K. (1973). *Counterfactuals*. Oxford: Basil Blackwell Press.
- Lewis, D. K. (1986). Counterfactuals and time's arrow, in *Philosophical Papers Volume 2*. New York: Oxford University Press. (Pp. 32-66).
- 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.
- McCawley, J. E. (1996). Conversational scorekeeping and the interpretation of conditional sentences, in M. Shibatani and S. A. Thompson, eds., *Grammatical constructions: Their form and meaning*. New York: Oxford University Press. (Pp. 77-102).
- Nozick, R. (1981). *Philosophical Explanations*. Harvard University Press.
- Oakley, I. T. (2001). A skeptic's reply to Lewisian contextualism, *Canadian Journal of Philosophy*, 31(3): 309-332.
- Sosa, E. (1999). How to defeat opposition to Moore, *Philosophical Perspectives*.
- Stanley, J. (2005). *Knowledge and practical interests*. New York: Oxford University Press.
- Williamson, T. (2002). *Knowledge and its Limits*. New York: Oxford University Press.

Williamson, T. (2005). Knowledge, context, and the agent's point of view. In Preyer & Peter, eds.: *Contextualism in Philosophy: Knowledge, Meaning, and Truth*. Oxford: Clarendon Press: 91-114.

Williamson, T. (2007). *The Philosophy of Philosophy*. Malden, MA: Blackwell Publishing.