

On Pragmatic Encroachment in Epistemology

Jeremy Fantl and Matthew McGrath

We argue, contrary to epistemological orthodoxy, that knowledge is not purely epistemic—that knowledge is not simply a matter of truth-related factors (evidence, reliability, etc.). We do this by arguing for a pragmatic condition on knowledge, KA: if a subject knows that *p*, then she is rational to act as if *p*. KA, together with fallibilism, entails that knowledge is not purely epistemic. We support KA by appealing to the role of knowledge-citations in defending and criticizing actions, and by giving a principled argument for KA, based on the inference rule KB: if a subject knows that *A* is the best thing she can do, she is rational to do *A*. In the second half of the paper, we consider and reject the two most promising objections to our case for KA, one based on the Gricean notion of conversational implicature and the other based on a contextualist maneuver.

According to received tradition in analytic epistemology, whether a true belief qualifies as knowledge depends only on purely epistemic factors – factors that are appropriately “truth-related.” If my true belief that *p* qualifies as knowledge while yours does not, this must be because of some difference in our evidence regarding *p*, the reliability of the processes involved in our beliefs that *p*, our counterfactual relations to the truth of *p*, and so on. My true belief cannot count as knowledge, and yours not, simply because you have more at stake than I do in whether *p*. Raising the stakes may indirectly affect whether one satisfies the belief condition on knowledge (because of one’s worrying about the costs of being wrong, for example), but it cannot otherwise make a difference to whether one knows; that is to say, stakes cannot affect whether one is in a *position to know*. If we think of purely epistemic factors as fixing a subject’s *strength of epistemic position* regarding a proposition *p*, we can formulate the traditional idea that knowledge is purely epistemic as a supervenience thesis, which we dub “epistemological purism”:

(purism) For any two possible subjects S and S', if S and S' are alike with respect to the strength of their epistemic position regarding a true proposition p, then S and S' are alike with respect to being in a position to know that p.

Some philosophers have questioned purism on the ground that the very idea of strength of epistemic position embodies a mistake (Williams 1992). Here we are interested in a different kind of dissent from the tradition, one which, in Jonathan Kvanvig's evocative phrase, argues for "pragmatic encroachment" in epistemology. According to Jason Stanley (forthcoming), John Hawthorne (2004), David Owens (2000), and Christopher Hookway (1990), a subject's stakes in whether p can be relevant to whether she is in a position to know. Our aim here is to argue for pragmatic encroachment by defending the following pragmatic (or better "practical") condition on knowledge:

(KA) S knows that p only if S is rational to act as if p.

'Acting as if p' is a familiar notion, but if an explanation is wanted, we may say that to act as if p is to do whatever one is rational to do on the assumption of p.¹ KA thus tells us that S knows that p only if S is rational to do whatever S is rational to do on the assumption of p.²

We claim that if KA is true, then the cost of purism is the rejection of another central plank in recent epistemological dogma, fallibilism:

(fallibilism) It is possible for a subject to have fallible knowledge.

¹ As a first approximation, S is rational to do A on the assumption of p iff, for no competitor B to A, is the outcome *p and S does B* rationally preferable to the outcome *p and S does A*.

² If we think of an act as rational iff it has the highest expected utility of the acts available to the subject, KA can be read as follows:

If S knows that p, then whatever act has highest expected utility for S on the assumption that p also has the highest expected utility for S in fact.

The act(s) with the highest expected utility on the assumption of p is (are) simply the act(s) that comes out with the highest expected utility conditionalizing on p.

In (Fantl and McGrath 2002), we use 'PC' for the following stronger thesis:

S is justified in believing that p only if S is rational to prefer as if p,

where S is rational to prefer as if p iff, for all states of affairs A and B, S is rational to prefer A to B on the assumption of p iff S is rational to prefer A to B, in fact. In the present paper, we replace talk of 'preferring as if p' with talk of 'acting as if p' because our purpose here is to respond to objections to our general approach, not to its details.

where S has fallible knowledge that p iff S knows that p but S's strength of epistemic position regarding p is not maximal with respect to justification (i.e., there are epistemic positions regarding p with respect to justification).³ If fallibilism is false, then its negation, infallibilism, is true. If infallibilism is true, skepticism looms, because for a great many of the mundane propositions we ordinarily take ourselves to know, our justification is non-maximal. You can know, in advance, that your spouse will be home around 5pm; but your justification improves once you see him walking in the door. You can know, sitting on your front porch, that there is a tree in your backyard, but your justification improves once you look out your back window. That one can improve one's justification in these ways can be confirmed by thinking of the bets one would rationally take before and after the improvement.

If KA is true then, barring infallibilism, purism is false. Suppose fallibilism and KA are true. Then there is a possible subject S1 who fallibly knows that p and so is rational to act as if p. Now, so long as a subject's epistemic position regarding p is non-maximal with respect to justification, her stakes in whether p can make a difference to whether she is rational to act as if p, independently of her strength of epistemic position regarding p and of the other traditional conditions on knowledge (e.g. belief, truth, proper basing, etc.). For example, if I find out that the police are (for the first time ever) about to ticket illegally parked cars on my quiet, rural street, my stakes in whether my car is legally parked rise. This makes a difference to whether I'm rational to act as if my car is legally parked, even if my strength of epistemic position doesn't change (I still remember—or seem to—legally parking it), my belief doesn't change (I'm just a confident guy), etc. Therefore, there is a possible subject S2 just like S1 with respect to

³ Some may object to our using 'fallibilism' in this way. Traditionally, the term is associated with a doctrine about evidence, viz. that one can know that p even if one's evidence is logically consistent with its not being the case that p. However, we think greater generality can be achieved by dropping the talk of evidence. Externalists are getting at a fallibilist idea when they insist that knowing that p doesn't require perfect reliability of the belief-process process, nor does it require having a belief that matches the truth across all possible worlds, etc.

strength of epistemic position regarding p and the other traditional conditions on knowledge, but who, because of her different stakes, is not rational to act as if p. By KA, S2 doesn't know that p. Therefore, because S2 satisfies all the other traditional conditions on knowledge (S2 believes that p, p is true, etc.), S2 must not be in a position to know that p. But S1 is. It follows that purism is false..⁴

Invariantist views that include KA fall under the category of “subject-sensitive invariantism.”⁵ In the first part of this paper, we will argue for KA. Our argument relies on the role knowledge-citations play in criticizing and defending actions. We then consider and reject what we see as the most promising attempts to resist our case for KA. The first proposes that KA, while false, encodes a conversational implicature associated with knowledge-citations. The second offers a contextualist response to our argument.⁶

Part I. The Case for a Pragmatic Condition on Knowledge

⁴ One might question the crucial assumption that stakes can affect rationality of acting as if p without affecting strength of epistemic position regarding p. We believe that if a subject S fallibly knows that p, then S is possibly such that S is offered a credible bet on whether p with the stakes required to make S not rational to act as if p, without affecting S's satisfaction of the traditional conditions on knowledge. If you are skeptical of this, then we can replace that assumption with the following weaker one: if it is possible for a subject to have fallible knowledge, then it is possible for a subject to have fallible knowledge of propositions on the truth of which she could possibly have had much more riding than she in fact does, so much so as to not be rational to act as if the proposition is true.

⁵ “Subject-sensitive invariantism” is DeRose's (2004) term. It is a broad term covering not only views, like ours, according to which there is a pragmatic condition on knowledge, but also views that make the non-salience of possibilities of error a condition for knowledge (Hawthorne 2004). One difficulty with this term is that it does not distinguish between purist and non-purist invariantism. Traditionally, purist theories of knowledge are subject sensitive as well – the only relevant factors to whether S knows that p are (epistemic) features of S's context. Better, perhaps, would be the label “pragmatic invariantism.” Nonetheless, “subject-sensitive invariantism” seems to be the currently favored option.

⁶ Some philosophers have recently suggested a role for salience of error in the conditions of knowledge (Hawthorne 2004, Hofweber 1999). Hawthorne (168-72), expresses guarded sympathy for the principle SC:

(SC) One cannot know that p if counterpossibilities of error are salient to one.

It is not obvious that SC is inconsistent with purism. Everything depends on how one thinks of the knowledge-destroying power of salience. Hawthorne suggests a number of models, including the “belief-removal” and the “evidential” models. If salience merely removes belief, it is not inconsistent with purism, which is, strictly speaking, a thesis about being in a position to know. If salience provides counterevidence, again, it is consistent with purism.

We often defend actions by citing knowledge. For example, I might say to my spouse, in defending my driving straight home without stopping off to get yams the night before Thanksgiving, “I know we have them at home.” In citing my knowledge, I am trying to convince my spouse that I am rational to act as if we have yams at home – that is, head home rather than stop off and buy some. Likewise, in defending an ambitious candidate’s decision to drop out of race, the refreshingly honest campaign adviser might say, “Well, he knew he was going to lose.”

We also criticize actions in an analogous fashion. Turning in her resignation, a tireless campaign staffer for Dennis Kucinich, who formerly was under the impression that Kucinich had a real shot at winning the nomination, might say to him, “Come on Dennis, you know you are going to lose. It’s time to quit.” A father might say to his daughter, Sue, who has missed work to make sure that his air conditioner is on, “Sue, what are you doing here? You knew I would turn it on. You should be at work.” Ordinary speakers feel that there is something amiss about continually checking the a/c, the bathroom sink, the kitchen stove. We say, “You knew everything was all right. Why did you keep checking?”, or “I knew they were locked; why did I bother going back?” Our favorite example of a knowledge-citing rebuke comes from the children’s book, *Owl Babies* in which a returning owl mother scolds her worried babies by saying, “What’s all the fuss? You knew I’d come back.”

So, knowledge-citations play an important role in defending and criticizing actions. Furthermore, this role is not restricted to subjects in low-stakes situations. When I say definitively, “I know we have yams in the house,” I aim to terminate a rather significant conversation about whether to stop at the grocer’s or not, and I say it even though I’m aware of the consequences if we don’t in fact have yams (Thanksgiving dinner without yams?!). Sue’s

father recognizes how much is at stake – it’s 100 degrees – yet confidently cites her knowledge in his rebuke. It’s precisely when the stakes are high that citing knowledge is such a powerful dialectical tool. This is not to say that stakes make no difference to our willingness to cite knowledge. But they make no difference to whether our knowledge-citations play the role we have described.

One might object that knowledge-citations play the role in question only because knowledge entails belief, and belief is fundamentally what plays this role.⁷ However, notice that when we defend or criticize actions by citing knowledge, it is irrelevant to us whether the subject actually believes or is psychologically sure of its truth or not. The staffer rebukes Kucinich for staying in the race, whether or not Kucinich believes that he’ll lose. If, in the course of a conversation, it becomes clear to the staffer that Kucinich does not believe he will lose, the staffer will certainly not withdraw his rebuke: “Oh, my mistake. I guess you’re right; you should stay in the race.” Rather, the reaction will likely be, “Look at the polls! You’re not even on the ballot in most states. You really should drop out.”

Likewise, when Dad scolds Sue for coming to check on the a/c, Dad knows that Sue isn’t sure and that she may not even believe that his a/c is on. He isn’t suggesting that she is acting contrary to her beliefs or what she is sure of, but rather that she is acting contrary to her evidence. True, Sue might say, by way of explaining herself, “I just wasn’t sure your a/c was on” but she knows this is not a real defense of her action. Her best defense is to try to convince

⁷ Similarly, one might object that knowledge citations play the role they do in defending and criticizing actions only because knowledge entails truth, and truth is fundamentally what plays this role. While we do sometimes cite the truth of some proposition to evaluate actions (“The winning numbers were 3, 7, 22, so it was good I chose those numbers” or “The winning numbers were 3, 7, 22, and you chose 3, 7, 21. Why couldn’t you have just picked 22?”), this has little to do with rationality of action. We don’t think that choosing 3, 7, and 22 was at the time more rational than any other relevant choice. If one’s strength of epistemic position regarding p is justificationaly insufficient, then although we may be glad we acted as if p, or regret that we didn’t, the mere fact that p is true does not make it the case that one is rational to act as if p. So, citing the fact of p should not convince us that a subject is rational to act as if p.

Dad that she did not have good enough evidence to know: “You sometimes forget to turn the a/c on, Dad; people die because of heat like this.” More to the point, she needs to convince her father that she wasn’t in a position to know. Our knowledge-citations work to defend or criticize actions, not because the success of those defenses or criticisms depend on whether the subject believes that p, but rather whether the subject is in a position to know that p.⁸

The fact that knowledge-citations play the role they do in defending and criticizing actions is evidence that KA is true. Consider again the yams example. When I defend my choice to drive home without stopping to buy yams, I cite my knowledge that we have yams at home as a reason for thinking my act is rational. And presumably, if I do know, my knowing is a *good* reason for thinking my act is rational. But how could the fact that I know that we have yams at home be a good reason for thinking my not stopping at the grocer’s is rational? The two claims seem unrelated. KA seems required to provide the needed link: “I know we have yams,

⁸ When we defend or criticize actions by citing knowledge, we are invoking a notion of rational action whose application is fixed by the subject’s total epistemic position (a function giving the subject’s strength of epistemic position for all propositions) and the value she attaches to various possible outcomes. This same notion of rationality is also at work in first- and second-person deliberation, singular or plural, about what to do. In deliberating, we regard the mere psychological fact of belief or its absence as of little relevance; what counts is one’s epistemic position regarding the relevant propositions. We take it that rational decision theorists aim to capture exactly this core notion with their apparatus of probabilities and utilities. We would only insist that the probabilities in terms of which the decision theory is developed reflect rational credences, or credibility values in Nozick’s (1993) sense, rather than purely psychological degrees of belief.

We do not deny the existence of more subjective senses of ‘rational’. When someone believes, and perhaps even justifiably believes, that she is rational to do A in the sense we have been describing, then even if she is wrong about this, if she does A, she can be said to be acting in accord with her own best judgment about what is rational; if she had acted otherwise, she would have been irrational by her own lights. This notion of ‘rationality/irrationality by one’s own lights’ is obviously important to our practices of excusing or at least understanding actions after the fact. But we think that this more subjective notion is parasitic on the core notion of rational action.

Arguably, there are also more objective offshoots of this core notion of rationality. One’s action can be criticized on the grounds that one ought to have known better. So, one might rebuke an overly hasty friend by saying, “You broke the VCR! You should have known that just pushing all the buttons would only make things worse.” The less our friend knows about electronic equipment the less happily we will call his action irrational (though we will still insist it is uninformed or unintelligent). It is worth noting that this “should have known” rebuke is further evidence for KA. Our friend should have known that pushing all the buttons on his VCR would make things worse. How does that function as a rebuke? Well, he should have known that, and had he known it, he would have been rational to act as if pushing the buttons would make things worse – that is, he would have been rational not to push the buttons. He deserves rebuke because he should have had knowledge that would make it rational to act in a way different than the way he did.

and so I am rational to act as if we do, and so am rational to just drive home and not stop by the grocer's.”⁹ Likewise, with Kucinich and his campaign staffer. In citing Kucinich's knowledge that he will lose, the staffer presents this as a reason for thinking Kucinich is rational to drop out. And KA seems needed if the reason is to be a good one: “Dennis, you know you're going to lose, and so you're rational to act as if you're going to lose, and so you're rational to drop out.” (Note that a second staffer might defend Kucinich against the first's rebuke by citing other motives to stay in the race: “It doesn't matter that he's going to lose. He's trying to get the message across.” Here, by citing this motive, the second staffer's reply blocks the reasoning from *Kucinich is rational to act as if he is going to lose* to *He should drop out*. Given his motives, the thought is, staying in the race is a way of acting as if one is going to lose. Without KA, it would be hard to see how the second staffer's reply exonerates Kucinich.) The truth of KA thus seems to be required if in defending and criticizing actions our knowledge-citations can provide good reasons for our judgments about actions.

This case for KA can be bolstered by a principled argument. Suppose I do know that we have yams at home, and suppose I also know that if we have yams at home, then driving home without stopping off at the grocer's is the best thing I can do – i.e. will have the best all-around consequences for me of all the acts available to me. If I have these two pieces of knowledge, then I also will know that driving home is the best thing I can do. But if I know *that*, then I am rational to drive home. More generally, it seems that the following argument is valid:

- 1) S knows that p.
- 2) S knows that if p, then A is the best thing she can do – that is, will have the best consequences of all the acts available to her.
- So, 3) S is rational to do A.

⁹ Another link is needed as well, of course—the link that specifies that driving home is what is rational to do, given that we have yams. Often, what it's rational to do, given p, is just obvious (as it is in the yams example). Sometimes, though, it's not, and this is what explains attempts to respond to knowledge-citing rebukes that involve the clarification of motives. See the above paragraph for such a response on behalf of Kucinich.

The statement of this inference's validity is the following principle:

(pre-KA) If S knows that p, then if S knows that if p, then A is the best thing she can do, then S is rational to do A.

In (Fantl and McGrath 2002), we argue that pre-KA can be strengthened to deliver KA. We will not rehearse that argument here. We assume, in other words, that if the argument (1) – (3) is valid, then KA is true.

The argument 1-3 consists of two more immediately compelling arguments:

Argument #1:

- i. S knows that p,
- ii. S knows that if p, then A is the best thing she can do,
- So, iii. S knows that A is the best thing she can do.

Argument #2

- i. S knows that A is the best thing she can do.
- So, ii. S is rational to do A.¹⁰

Argument #1 follows from the closure of knowledge under modus ponens.¹¹ Argument #2 has no standard name. We will call it the “knowledge-to-action” inference, or KB.¹² Much of this paper will be devoted to answering challenges to KB.

At the expense of some elegance, one might argue against fallibilist purism using pre-KA rather than KA. The argument requires two premises. The first is that if it is possible to have

¹⁰ Arguments #1 and #2, and so our inference 1-3 above, hold for ‘being in a position to know’ as well as knowledge. We neglect this distinction except when something particularly important hangs on it.

¹¹ As we note in (Fantl and McGrath 2002), some instances of closure like Argument #1 are peculiar. For example, if A’s having the best consequences depends causally, logically, or perhaps even probabilistically on whether p, the argument is a bit fishy. Suppose Bob wants to get to Providence by train asap. The conductor of the train, a sworn enemy of Bob’s, informs Bob that if Bob boards the train, he will not take it to Providence, but if Bob doesn’t he will take it to Providence. Now suppose Bob decides not to take the train. Then it seems Bob knows it will go to Providence, and so knows the best thing to do is to board it (because he wants to go to Providence). So Bob is rational to board it. But let us put aside such peculiar cases.

¹² In appealing to Arguments #1 and #2, we separate ourselves from Jason Stanley (forthcoming) on the question of whether it is actual or perceived stakes that make an epistemic difference. The second premise of Argument #1 makes it clear that we are not committed to claiming that stakes that are beyond one’s ken can make an epistemic difference. We do, however, think that reasonable but false beliefs about stakes can also make an epistemic difference.

fallible knowledge at all, then it is possible for there to be a subject S1 and a proposition p such that (i) S1 fallibly knows that p and (ii), S1 knows that if p is true, then A is the best thing S1 can do, where A is the “efficient” act and B the “safe” act (e.g., A might be going directly home and B stopping by the grocer’s).¹³ The intuitive thought behind the first premise is simply that if fallible knowledge is possible at all, then it is surely possible to have fallible knowledge of the premises of a modus ponens argument, the conclusion of which is that a certain efficient act is the best thing one can do. The second premise is that, so long as a subject’s epistemic position regarding p is justificational non-maximal, the costs associated with $\sim p$ can make a difference to whether she is rational to perform the efficient act A or the safe act B, independently of her strength of epistemic position regarding p, the other traditional conditions on knowledge (e.g. belief, truth, proper basing, etc.), and her knowing that if p then A is the best thing she can do.

Suppose fallibilism and KA are true. Given the first premise, there is an S1, who knows that p and also knows that if p then A is the best thing she can do. Given pre-KA, S1 is rational to do A. Given the second premise, there is a possible subject S2 just like S1 except that S2, because of her different stakes, is rational to do the safe act B rather than A. By pre-KA, S2 does not know that p and so is not in a position to know that p. But S1 is. Thus, purism is false.

It is therefore not essential to our overall argument that pre-KA can be strengthened to deliver KA. Nonetheless, we think that KA is needed if knowledge is to play the role it does in defending and criticizing action. Pre-KA applies only to cases in which the subject both knows that p and knows that, if p, then some act A is the best thing she can do. In the majority of the above examples, this is precisely the situation the subjects are in. I know that if we have yams at home, driving straight home is the best thing I can do. Sue knows that, if her father’s a/c is

¹³ Condition ii is true in every case in which we know that, if some proposition is true, the efficient act will have the best consequences. So, for example, I know that, if the gas is off, not checking will have the best consequences; and I know that if class is cancelled, staying home will have the best consequences; etc.

already on, going to work is the best thing she can do. But many cases – perhaps most – are not like this. In many cases in which we know that p , we won't know about any particular course of action, that if p , then that action will have the best consequences. Suppose I know that the forecast is for rain. This isn't usually enough to know that it will rain and suppose, in this case, I don't. Then I may well *not* know that if the forecast is for rain, then bringing my umbrella is the best thing I can do. For all I know, it won't rain, and if it doesn't, bringing my umbrella will have been a slightly inconvenient and useless exercise. Despite this, I may still be rational to act as if rain has been forecast – and so to bring my umbrella. And I can be rebuked in this case for not acting in the required way: “Why didn't you bring your umbrella? You knew rain was in the forecast!”

Pre-KA is silent on these kinds of cases, but KA is not. KA, unlike pre-KA, has the implication that, if I know that rain has been forecast, then I am rational to act as if it has been forecast – that is, bring my umbrella. Therefore, though pre-KA is adequate to argue against epistemological purism, the more comprehensive principle is the simpler and more elegant KA.

One might worry that, though the argument for KA is strong, KA cannot be right because it has skeptical or at least semi-skeptical implications. Hawthorne, for example, claims that he “would happily bet a penny against a billion dollars on not having been born in England, and on the falsity of the law of noncontradiction” (2004 176 ff.). If Hawthorne's cheer at such possible bets indicates that he would be rational to bet the penny, and if to bet the penny is to fail to act as if he was born in England or as if the law of noncontradiction is true, then it seems to follow, given KA, that if Hawthorne were offered these bets, he would not know he was born in England and would not know the law of noncontradiction. (We could get fussy about the value of a

penny, but we'll spare the reader! Let's assume that other things being equal a penny is better kept than thrown out.) But surely he does know these things! Thus, KA must be wrong.

Notice that the skeptical consequences do not follow merely from KA, but from KA together with purism. Hawthorne is in fact being offered no such bet. He in fact is rational to act as if he was born in England because, of all his *available* options, what he's rational to do is also what he's rational to do on the assumption that he was born in England. Therefore, KA is consistent with the claim that he knows. Offered the bet, however, he would not be rational to act as if he was born in England, and so, by KA, would not know. But it doesn't follow that as things in fact stand he doesn't know, unless we suppose that sameness of epistemic position ensures sameness with respect to being in a position to know. (The offering of the bet would not affect his strength of epistemic position regarding these propositions.)

So KA by itself does not have skeptical implications. But if one rejects skepticism (and the associated doctrine of infallibilism) and accepts KA, it seems one will be committed to some rather counterintuitive claims. One will have to say that being offered credible bets may strip one of knowledge. Doesn't that sound just as ridiculous as the claim that one knows that p but isn't rational to act as if p? If it does, then it looks like fallibilism, KA and purism jointly constitute a paradox. Each seems true, yet they are jointly inconsistent.

We claim that if infallibilism is out of the question – and we assume it is – then there is better reason to deny purism than KA. KA follows from an apparently innocent use of closure and the extremely intuitive principle, KB: if you know that A is the best thing you can do, you are rational to do A. It is this argument for KA, we think, that gives it an edge over purism. Return to the hypothetical bet. One might find a way to reassure oneself that, odd though it sounds, if Hawthorne were offered the bet, he would still know that he was born in England but

wouldn't be rational act as if he was, i.e., refuse the bet. But it is much harder to reassure oneself that in this situation Hawthorne would know that refusing the bet is the best thing he can do and yet be rational to refuse it. That sounds absurd, and not because of a definitional cheat. We do not define 'the best thing one can do' as 'the rational thing for one to do'. A thing's being the best thing one can do is a matter of what its actual utility is, not its expected utility. But actual and expected utility are related in the manner KB suggests: if you know that A will have the highest actual utility of the options available to you, then A has the highest expected utility of the options available to you.

In the remainder of this paper, we will consider attempts to resist our case for KA. To resist our case requires not only giving an alternative explanation of the role of knowledge-citations in defending and criticizing actions but also explaining how the principled argument for KA invoking KB goes wrong.

Part II. Two Attempts to Resist the Case for KA

1. *Conversational Implicature*

One alternative for the purist is to argue that, while KA is false, when we cite knowledge that p in defending and criticizing actions, we conversationally implicate that the subject's epistemic position is strong enough to make her rational to act as if p. And this is what explains how knowledge-citations can be used to support judgments about actions. It is not that we appeal to a subject's *knowing* that p as a reason for the relevant judgment about her action. Rather, by citing this knowledge we imply that the subject's epistemic position is strong enough to make her rational to act as if p, and it is the latter claim that is the reason for our judgment about the subject's action. For example, suppose your husband suggests that you go back to

check that the gas is turned off on the stove before you leave for a trip. And suppose you defend your not going back by saying, “Don’t worry, I know the gas is turned off.” By saying this – the thought is – you imply that your epistemic position is strong enough to make it rational for you to act as if the gas is turned off – that is, continue down the road rather than turn back to check. This KA-implication¹⁴ is only conversational, not semantic, however, so it doesn’t tell us anything about the truth conditions of knowledge-citations.

In Grice’s (1967 32) familiar petrol case, a stranger, standing by his “clearly immobilized car,” tells you “I am out of petrol.” If you say in response, “There’s a garage around the corner,” you conversationally implicate that the stranger can purchase petrol there. On the assumption that you are a cooperative speaker, i.e., that you are obeying what Grice calls the “Cooperative Principle,”¹⁵ you would say, “There’s a garage around the corner” only if you intend your addressee to conclude from your saying this in the circumstances (and from your cooperativeness) that he can buy petrol there. In a situation like this in which you knew the garage was closed, the stranger later could rightly accuse you of lying, of knowingly leading him to a false belief.¹⁶

Under the current proposal, then, when we cite S’s knowledge that p to defend or criticize S’s actions, the assumption that we are obeying the Cooperative Principle will enable our interlocutors to conclude that S has an epistemic position strong enough to make her rational to

¹⁴ We mean “implication” to be neutral between semantic entailments and conversational implicatures.

¹⁵ The Cooperative Principle is this: “Make your conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged” (Grice 1967 26).

¹⁶ Grice defines conversational implicature as follows: “A man who, by (in, when) saying (or making as if to say) that p has implicated that q, may be said to have conversationally implicated that q, provided that (1) he is to be presumed to be observing the conversational maxims, or at least the Cooperative Principle; (2) the supposition that he is aware that, or thinks that, q is required in order to make his saying or making as if to say p (or doing so in *those* terms) consistent with this presumption; and (3) the speaker thinks (and would expect the hearer to think that the speaker thinks) that it is within the competence of the hearer to work out, or grasp intuitively, that the supposition mentioned in (2) is required” (Grice 1967 30-1).

act as if p . If your husband tells you he knows the stove is off, as a way of defending his not going back to check, your reasoning would be explicitly represented thus: “He wouldn’t have said he knows, given this situation, unless his basis for believing the gas is off were good enough to make him rational to act as if it is, and so to make him rational to continue down the road.”

We want briefly to raise two problems for this proposal, and then examine a third problem in some detail. The first is that this proposal cannot easily explain the attractiveness of instances of KA or of KA itself. When one implicates that x is G by saying that x is F , generally the conditional ‘whenever something is F it is G ’ does not seem true. Thus, ‘whenever there is a garage around the corner, one can buy petrol there’ seems clearly false. However, the corresponding conditional, ‘whenever a subject knows he has yams, the subject is rational to act as if he does’ does not seem clearly false. Rather, it seems true. And similar considerations hold not just for instances of KA, but for KA itself. The second problem is that the proposal by itself seems unable to explain why we often reason in conformity with the contrapositive of KA. After an argument over the spelling of a word, in which you claimed you knew how it was spelled, I might catch you looking the word up in a dictionary: “Well, if you knew how it was spelled, you wouldn’t have had to check, would you?! You don’t know!” Similarly, if I’m unwilling to take your bet on whether p , you might say to me “So, you don’t know that p , do you?!” Here, if there is implicature, it is not by courtesy of our citing knowledge.

A third problem is even more pressing. If KA-implications were merely conversational, one would expect there to be many cases in which one can respond to a knowledge-citing rebuke by saying “Well, I do know that p , but I’m not rational to act as if p .” That is, one would expect KA-implications to be cancelable, and in fact, to be regularly canceled. But, are they?

Grice took conversational implicatures to be cancelable, because a speaker can always make it clear that she is opting not to abide by the Cooperative Principle, at least at non-fundamental levels. So, saying “There is a garage around the corner, but you can’t buy petrol there” nullifies or cancels an otherwise clear implication. An implication *that q* of saying “p” is cancelable if it is admissible to add “but not q” to one’s claim “p” (Grice 1989 44). Here we will read ‘admissible’ narrowly. The question is whether one can add “but not p” without absurdity, not whether adding this would be *in some way* odd (e.g., unmotivated, smug, cute, or precious) in the conversational context. Typically, cancelations will be rather odd. Your interlocutor in the petrol case would be perplexed by your saying, “There’s a garage around the corner but you can’t buy petrol there.” However, by saying this you clearly cancel the implication, and in fact may speak the truth.¹⁷

Are KA-implications cancelable? Suppose that, by saying “Sue knows her father’s a/c is on,” you conversationally implicate that Sue’s epistemic position is good enough to make her rational act as if her father’s a/c is on – that is, go to work. You should then be able to specify without absurdity that Sue is not rational to act as if the a/c is on. But can you? It does not sound altogether happy to say, “Sue knows that the a/c is on, but she is not rational to act as if it is on,” and its unhappiness seems more a matter of absurdity than it does quaintness or humor. But, admittedly, it is not clearly absurd in the way “Sue knows that p, but not-p” is.

We do not find this unclarity particularly surprising. Semantic implications are not in every case obvious, and the less obvious they are the less clear it will be whether they are cancelable. Consider the implication from (p) *My mother’s brother’s father’s only grandchild is 34* to (q) *I am 34*. The latter is a semantic implication of the former, but one can append “but not

¹⁷ Similarly, it is hard to see how a woman could say, “I had a romantic date with a man, and in fact it was my husband,” without meaning to be humorous. But, obviously, canceling a conversational implicature can be a source of humor.

q” to a statement of “p” without obvious absurdity: “My mother’s brother’s father’s only grandchild is 34, but I am not 34.” The difficulty derives from the fact that semantic absurdities can fail to be obvious, and so can be canceled, even while they retain their semantic status. We do think that our habit of citing knowledge to defend or criticize actions is evidence for KA, and we claim that KA is intuitively attractive, but we do not claim that KA is so obvious that the implications it encodes cannot be canceled without obvious absurdity.

However, it’s not enough for the opponent of KA to show that we can cancel implications of “S knows that p.” The opponent also has to show that we can cancel implications of “S knows that A is the best thing she can do.” Recall that our principled argument for KA can be broken down into two steps – an innocent use of closure, and a knowledge-to-action inference, KB, according to which if S knows that A is the best thing she can do then S is rational to do A. Given these two steps, KA follows. We continue to assume that closure is valid. Therefore, KA cannot be conversational unless KB is invalid. But then we can ask why KB seems so clearly valid. In the context of the current objection, it looks like our opponent has a ready answer. If KA admits of a conversational treatment, KB does as well. When we cite S’s knowing that A is the best thing she can do in defending or criticizing S’s action, this is just another case of conversationally implicating that S has a strong enough epistemic position regarding p to make S rational to act as if p. After all, to act as if A is the best thing one can do is exactly to do A.¹⁸

So, the central question is whether KB-implications are cancelable. If they are not, they are not conversational implicatures, and if they are not conversational implicatures, then neither are KA-implications. Moreover, if KB-implications *are* cancelable, then because KB is obvious (unlike the implication in our example about “one’s mother’s father’s only grandson”),

¹⁸ That’s because to act as if p is to do what it is rational to do on the assumption of p. On the assumption that A is the best thing one can do, the thing it’s rational to do is A. In decision theoretic terms, on the assumption that A will have the highest actual utility, A has the highest expected utility.

cancelability is good evidence for KB's invalidity. The remainder of this section will therefore focus on the question of whether KB-implications are cancelable.

Consider an attempt to cancel a KB-implication:

(1) Sue knows that going to work is the best thing she can do, but she is not rational to go to work.

This sounds absurd, as does an attempt to cancel the related implications KB1 and KB2

KB1:

- i. S knows that (act) A is worse than (act) B.
- So, ii. S is not rational to do A.¹⁹

(2) Sue knows that checking on the a/c is worse than going to work, but she is rational to check on the a/c.

KB2:

- i. S knows that A is the worst thing she can do.
- So, ii. S is not rational to do A.

(3) Sue knows that checking the a/c is the worst thing she can do, but she is rational to check on the a/c.

1, 2, and 3 are absurd. We believe that the natural explanation for their absurdity is that they state impossibilities, but all we are arguing here is that the implications in question are not cancelable.

¹⁹ One might argue against fallibilist purism simply from closure and KB1. Closure and KB1 entail something very close to pre-KA, viz. that if S knows that p then, if S knows that if p then A is worse than B, then S is not rational to do A. The argument against fallibilist purism would be exactly analogous to our earlier argument using pre-KA.

One might reply that there are ways to cancel implications other than by simply adding “but not q” to the statement “p”. Additional explanatory material may be supplied. The following are attempts to cancel KB, KB1 and KB2 (hereafter lumped together into “KB-implications”) by pointing out that, while the subject might know the relevant proposition, the subject doesn’t know “for sure”:

(1*) Sue knows that going to work is the best thing she can do, but she is not rational to go to work because she doesn’t know this for sure, and if she’s wrong, the consequences will be disastrous.

(2*) Sue knows that going to work is worse than checking the a/c, but she is rational to go to work because she doesn’t know this for sure, and if she’s wrong, the consequences will be disastrous.

(3*) Sue knows that going to work is the worst thing she can do, but she is rational to go to work because she doesn’t know this for sure, and if she’s wrong, the consequences will be disastrous,

We think 1*, 2*, and 3*, and variations on them, do not succeed, for two reasons.

First, we suspect that in uttering 1*, 2*, and 3* one retracts, or hints at retracting, one’s original knowledge attribution. This suspicion can be confirmed by considering the effect of repeating the knowledge attribution at the end of these statements. Were there no retraction or suggestion of retraction in 1*, 2*, or 3*, we should expect that a speaker would not seem to vacillate in repeating the knowledge attribution. In the petrol case, if you say, “There is a garage around the corner, but they are out of petrol; nevertheless, there is a garage around the corner,” you do not vacillate – you are simply being redundant. By contrast, when you repeat the knowledge attribution at the end of 1*, 2*, or 3*, you do seem to vacillate on the question of whether Sue knows.²⁰

²⁰ One possible explanation for this vacillation is that, after asserting something that clearly implies that p, asserting that if not p, then q pragmatically implies that one does not know that p. For discussion of purported implicatures associated with conditionals see Levinson (1983 213).

The second consideration is more decisive. Note that the apparent absurdity of a statement will tend to diminish with greater length and complexity. The demonstrable absurdity of 4, for example, may not be immediately obvious:

(4) Sue knows that going to work is the best thing she can do, but she is not rational to go to work because she doesn't know this for sure, and, at any rate, she doesn't know that going to work isn't worse than checking on her father's a/c, just to be safe.

But 4 entails something absurd – that one can know that an act is best but not know that it isn't worse than some other act. Like 4, 1*, 2*, and 3* have entailments that are clearly absurd,

(5) Sue is not rational to do something she knows is the best thing she can do.

(6) Sue is rational to do something she knows is worse than another given option.

(7) Sue is rational to do something she knows is the worst thing she can do.

Compare these with the paradigmatically successful cancelation in the petrol case: “There is a garage around the corner, but you can't buy petrol there.” This cancelation has no peculiar entailments at all, let alone absurd ones. We take this difference to cast serious doubt on the cancelability of KB-implications, whether by using the language of “knowing for sure” or by any other strategy. The absurdity of what all such attempted cancelations entail renders those attempts failures, however intuitively appealing the attempted cancelation may be.²¹

We conclude that our case for KA cannot be resisted by treating KA-implications as conversational implicatures. There may be a way to cancel KA-implications, or at least make

²¹ An adequate cancelation test must make room for the relevance of entailments. How we make room for entailments depends on how we think of role of *absurdity* in determining the results of a cancelability test. Is it genuine absurdity (apparent or not) that determines the test's results, or apparent absurdity? It makes little difference in the end. We have been assuming that it is genuine absurdity that matters. If so, then because having absurd entailments suffices for being absurd, facts about entailments are relevant to the question of whether an attempted cancelation succeeds. But others may prefer to understand the cancelability test in terms of apparent absurdity. If apparent absurdity is what matters, we would need to reconstrue *cancelation* in a way that makes a role for entailments. We could say that an attempted cancelation succeeds iff neither it nor any of its entailments of it appears absurd. Given this understanding of cancelation, 1*, 2*, and 3* will count as failed attempts. In fact, we think that attempts at canceling KA-implications will also count as failures on this understanding of cancelation, assuming the validity of our argument from KB, through closure, to KA.

We should make clear that we are not denying that apparent absurdity (lack of it) is good *prima facie* evidence for uncancelability (cancelability). We are insisting only that it is not the only relevant evidence.

cancelations of KA-implications not seem as obviously absurd as some cancelations. But the real issue concerns KB-implications. These, we have argued, are simply uncancelable and so not conversational implicatures.²² And if they are not conversational implicatures, we need to be given some explanation of their apparent force other than the natural one according to which they are genuinely valid implications.

2. *Contextualism*

Here is an initially attractive explanation of the role of knowledge-citations in defending and criticizing actions: citing knowledge is not, *contra* the implicature objection, merely a way of drawing our interlocutor's attention to our real basis for judging an action. In citing knowledge, we present knowledge that p itself as a reason in favor of or against acting as if p. But knowledge that p can be a reason here without *entailing* that the subject is rational to act as if p. A basketball coach may cite an athlete's tallness in defending the choice to play her as center. This doesn't require that there be some entailment between tallness and being suitable to play center, but only that this athlete's tallness, given other relevant information in the case, provides good reason to make her center. Similarly, when I say "I know we have yams in the house, and so I'm not going to stop by the grocer's," I present my knowledge as a good reason, relative to certain other information, including information about stakes. My knowledge, together with this other information, is sufficient reason to think that I am rational to act as if we have yams in the house, and so not to stop by the grocer's.

²² We are not denying that attributing knowledge often involves conversational implicature or other sorts of pragmatic implications. Perhaps, when one says, "I know that p," one implies that one is ready to justify one's belief in p, or that one has some evidence for p that the hearer does not have, and so on. We do deny, though, that KB- and KA-implications are conversational implicatures. For a more detailed argument that KB-implications are not plausibly pragmatic implications of any sort (e.g., speech-act implications, conventional implicatures), see the Appendix at the end of this paper.

This is a very appealing explanation. However, there is a stumbling block. Given the right situation, an assistant coach might reasonably say, “Yes, she is tall; but she is a bad choice for center – she can’t jump very high.” Likewise, one would expect – if knowing that *p* doesn’t entail being rational to act as if *p* – that given the right situation one could reasonably say “Yes, he knows we have yams in the house, but he isn’t rational to act as if we have yams in the house – too much is at stake.” And similarly for KB: “Yes, he knows that driving directly home is the best thing he can do, but he isn’t rational to do so – too much is at stake.” That is to say, the stumbling block is the apparent absurdity of conjunctions violating KA and KB.

Here is where a contextualist maneuver may help. One might concede that the attempted cancelations of KA and KB, 8 and 9, fail:

(8) He knows that we have yams at home, but he isn’t rational to act as if we do.

(9) He knows that driving straight home is the best thing he can do, but he isn’t rational to drive straight home.

and that they fail because they are plainly false. But perhaps they are only false (let alone plainly so) when we attend to them (or take them seriously, etc.).

How could this work? Stewart Cohen (manuscript) has argued that “*S* knows that *p*, but can’t use *p* in practical reasoning” is sometimes true but never assertable, because raising the issue of using *p* in practical reasoning can make certain possibilities of error salient – if the stakes are high enough – thus rendering the knowledge claim false. Cohen is responding to John Hawthorne (2004), but his response could be applied equally well to KA and KB. Violations of these principles can be true but raising the question of their truth tends to make them false because certain possibilities of error become salient. This may sound like a familiar contextualist move. For example, the contextualist allows that in ordinary contexts, “I know I’m

not a brain in a vat” is true, but raising the issue can change the context, driving up the standards for knowledge so as to make it false.

Contextualism might seem to be the best way for the epistemological purist to reply to our argument. As we have seen, it seems to remove the main stumbling block for a purist explanation of the role of knowledge-citations in defending and criticizing actions. But it also seems to provide the purist with a way to explain how our principled argument for KA through KB goes wrong. As with KA, KB-violations are unassertable, because bringing them to mind tends to make salient certain possibilities of error, which in turn drive up the standards for knowledge high enough to make those violations plainly false. So, KB is invalid, but true violations of it are unassertable.

There are at least four difficulties with this contextualist account.

The first is a difficulty for the claim that KA- and KB-violations are unassertable. Salience, as Cohen notes, can be overridden. We sometimes truly say, “Give me a break, of course I know my lottery ticket will lose.” So, if salience is responsible for the apparent falsity of violations of KA and KB, we should, in certain circumstances, be able to override it. We should be able to make (8) and (9) assertable. But they seem never to be assertable.²³

The second difficulty is that any adequate explanation of how our principled argument goes wrong must provide some account of why KA seems true and, especially, KB seems valid. This means doing more than just explaining why we cannot assert violations of them. Why can we positively assert the principles themselves?

²³ Keith DeRose (manuscript) claims that, when we are thinking about whether a third party subject may properly use a premise in practical reasoning (or to assert it), we “should” use the standards for knowledge operative in the subject’s own context. It is not clear to us that this is an appeal to salience. If it is, then it doesn’t avoid the overridability problem. If it isn’t, we don’t understand why we *should* change our standards like this, unless KA is true.

One simple explanation for the assertability of these principles is that attending to them raises our standards to skeptical heights, thereby making false all of the antecedents of all of their instances and so making them true. This is extremely implausible. For one thing, attending to the principles does not make false all of the antecedents of *all* of their instances, because even after attending to the principles, we are still more than willing to cite knowledge to defend or criticize actions. In fact, were the standards for knowledge to rise to skeptical heights, we should expect that, when ordinary folk accept KA or KB, they are somewhat willing to deny themselves knowledge that Bush is president, and even that they have hands. Of course, they are do not. Finally, it is unclear how our standards would rise. When we attend to the principles, what possibilities of error become salient and for which propositions? And, even assuming that the standards do rise to skeptical heights, we should be able to detect the presence of these standards.

A second explanation is that whenever we consider an instance of these principles, that instance becomes plainly true (because our standards rise sufficiently). We then mistakenly generalize to the principle itself, which remains false, whether we attend to it or not. This explanation is problematic. First, it is psychologically unrealistic: when we are deciding whether KA (or KB) is true (valid), we needn't consider instances at all, let alone make some sort of generalization from them. When thinking about KB, for example, we just find it intuitively compelling that if someone knows that A is the best thing she can do, then she is rational to do A. Perhaps the use of a variable 'A' indicates that when we think about KB in the abstract we are really thinking of an "arbitrary" instance. But it is not clear how the salience account could be applied to arbitrary instances. Second, it seems ad hoc to explain the apparent truth of KA and its instances in such radically different ways. The instances of KA, we are told, seem true because they are true (when we attend to them). KA itself, on the other hand, seems true because

of a mistaken generalization. Better, we think, would be an explanation according to which KA and its instances have the same standing.

A third problem for the contextualist objection is that it is out of keeping with standard contextualist methodology. Contextualist solutions to epistemological problems typically are designed to accommodate the context-invariant truth of apparent first-order axioms figuring in those problems. The typical contextualist diagnosis is that, in applying one or more first-order epistemic axioms to arrive at a pessimistic knowledge claim “S doesn’t know that p,” after beginning with its optimistic negation, “S knows that p,” we ignore a contextual shift in standards for knowledge; so, although we think we have contradicted our earlier claim, in fact we haven’t. This methodology is clear, for example, in the familiar skeptical application of closure, which seems to lead us to contradict the belief that we know we have hands:

“I know I have hands, obviously. But, wait, I don’t know I’m not a BIV. But if I don’t know *that*, then how can I know I have hands? So, I guess I don’t know I have hands.”

The familiar contextualist diagnosis is then this: in the ordinary context in which we begin, “I know I have hands” is true and “I don’t know that I’m not a BIV” is false. But when the possibility of being a BIV becomes salient, the standards for knowledge rise so as to make the latter claim true. We then use closure to conclude “I don’t know I have hands,” seemingly contradicting what we said or thought earlier.

Why doesn’t the contextualist block skepticism instead by denying closure?

Contextualists typically reply by appealing to the absurdity of conjunctions like, “I know I have hands but I don’t know I’m not a BIV,” suggesting that the absurdity is best explained by the hypothesis that closure is invariantly valid. But this is exactly *not* what our committed contextualist says about KA or KB. Instead, she says that these principles are false, despite the seeming absurdity of their violations. Notice that salience is *not* used to explain the

unassertability of violations of closure. As we noted earlier, if salience were used, we ought to be able to override it to assert such violations. We can't, and so contextualists take this to be reason to think closure is invariantly valid. But violations of KA, and especially KB, are always unassertable, too. Nor, finally, does the contextualist argue that we mistakenly generalize from the truth of instances of closure to the general principle. Instead, she simply accepts the principle. All this raises the uncomfortable question: why treat closure and the principles KA and KB so differently, when the very same reasons for taking closure to be valid apply to KA and KB as well?

The fourth difficulty is that contextualism does not take a subject's stakes into account when the knowledge-attributor is not aware of them. Cohen himself notes a related difficulty in contextualist solutions to the Gettier problem. The Gettier problem, Cohen points out, is the problem of "giving a general account of why it is that certain cases of justified true belief fall short of knowledge" (1998 519). According to David Lewis' contextualist solution, when a speaker thinks that there are Gettier-like features in a subject's situation, certain counterpossibilities resembling actuality become salient to the speaker, thus making "S doesn't know that p" true in the speaker's mouth (Lewis 1996). So, to use Cohen's example, if the speaker thinks that a subject, S, is looking at a sheep-shaped rock that occludes a sheep behind it, then possibilities of error in which there is no occluded sheep become salient to the speaker, making "S doesn't know that there is a sheep on the hill" true.

Cohen argues that when a third party speaker wrongly believes a subject isn't in a Gettier case, the speaker doesn't speak truly in saying "S knows." We have a stable intuition that when a subject is in such a Gettier situation, the subject simply doesn't know and no speaker can speak truly in attributing knowledge to him. Thus, the salience to the speaker of certain Gettier-style

counterpossibilities of error cannot be what makes it true to say “S doesn’t know” when in fact the speaker understands the subject’s situation correctly. What is salient to the speaker is irrelevant to whether “S knows” is true in the speaker’s mouth.

Here’s how this applies to the epistemic relevance of stakes: suppose that Mary and John, in Cohen’s (1999) airport case, have a great deal at stake in whether the plane stops in Chicago, though the speaker, Smith, believes that they, like him, have nothing much at stake. Smith, reflecting on Mary and John, won’t think of unusual counterpossibilities of error. For the contextualist, Smith’s reflection will not then raise the standards associated with ‘know’ in Smith’s mouth. What affects the standards contextually associated with ‘know,’ it seems, cannot depend on subject-determined factors beyond the speaker’s ken, but only on the speaker’s beliefs, doubts, etc. as well as what has been said in the conversation. So, it seems that the fact that Mary and John are in a high-stakes situation will not raise the standards associated with ‘know’ for Smith. So, he will speak truly in saying, “Mary and John know that the plane stops in Chicago.” But this seems wrong. We have a stable intuition – one that Cohen brought out clearly in his original example – that, given Mary and John’s high stakes, and the fact that they rationally should check further, they don’t know that the plane stops in Chicago. As with the Gettier case, we think that neither Smith nor anyone else can speak the truth in attributing knowledge to Mary and John. This intuition is reason to think that the epistemic relevance of stakes cannot be accounted for in a contextualist fashion. This argument is even clearer with KB. We have a strong intuition that, given that Mary and John are rational to check further, they do not know that boarding without checking is the best thing they can do. This is reason to think that misinformed Smith speaks falsely in attributing them knowledge.

Given the difficulties of the contextualist attempt to resist the case for KA, one might wonder why contextualists shouldn't just grant its truth. The main difficulty concerns the relation between contextualism and purism. As we noted earlier, the conjunction of KA and purism entails infallibilism, and so brings with it the threat of skepticism. No contextualist maneuver will elude the consequence. The contextualist therefore faces the same choice the fallibilist invariantist does, whether to respect the intuitively compelling principle of purism and reject the intuitively compelling principle KA, or vice versa. She cannot respect both. But it is no coincidence that contextualists are disposed to reject principles like KA. The main appeal of contextualism is the promise of preserving purism. The union of contextualism and purism enables the contextualist to vindicate our tendency to apply to third-party subjects the same standards we apply to ourselves. High-stakes Mary and John can truly say about low-stakes Smith, "We don't know, and neither does he."²⁴ Such a remark will be false under fallibilist subject-sensitive invariantism. To abandon purism is to abandon the main advantage of contextualism over subject-sensitive invariantism.²⁵ Contextualists must therefore resist our case for KA. But, as we have shown, the contextualist response to our arguments is unsatisfactory.²⁶

²⁴ It is unclear how much of an advantage purism is to contextualism, in the end. If Smith applies his standards to Mary and John, we get the result that "Mary and John know the plane stops in Chicago" is true in Smith's mouth. This seems wrong, since they are right to check further, to withhold belief. Conversely, although speakers in Mary and John's position will tend to deny Smith knowledge, this tendency may not be worth vindicating. Smith needn't check further. He has good evidence. He is right to have made up his mind. Is it so clear that Mary and John are correct in denying him knowledge?

²⁵ It may be that the main appeal of contextualism lies not in its consistency with purism, but in its ability to make sense of certain cases – namely ones in which our intuitions lean in favor of saying that two speakers in different contexts both speak truly when one says, "She knows that p" and the other says, "She doesn't know that p." If that is the case, then a contextualist can do as we suggest and reject purism while retaining her contextualism. If she accepts our arguments, she may also accept PC. That is, she will accept both that the content of knowledge attributions can vary with differences in speaker context and that PC is invariantly true. The resulting view we might call "subject-sensitive contextualism." Nothing we say excludes the possibility that such a view is true. In fact, depending on whether there are intuitive examples that favor a contextualist treatment, it might be a view we would endorse.

²⁶ We would like to thank an audience at the 2004 Kline conference at the University of Missouri-Columbia, as well as Tom Crisp, Anthony Everett, Richard Fumerton, John Hawthorne, Jon Kvanvig, Joe Salerno, Jason Stanley, and especially Baron Reed and two anonymous referees for their helpful comments, criticisms, and suggestions.

IV. Conclusion

Suppose you are convinced by the claim that knowledge-citations play a role in defending and criticizing actions and corresponding principled argument, through KB, for KA. But you may still find yourself pulled by your purist intuitions. As a way to remove the conflict, you might be tempted, then, to think, “Well, there are just different senses of ‘knowledge’. In order to know₁ one must be rational to act as if what is known₁ is true. The strength of epistemic position required for knowledge₁ will slide up and down the purely epistemic dimension, depending on stakes. (Some might claim that *knowing for sure* is knowledge₁, though one need not.) But there is another – purist – sense of ‘knowledge’ (knowledge₂) and the strength of epistemic position required for it will occupy a fixed threshold on the purely epistemic dimension. Our ordinary language term ‘knowledge’, you might think, is just ambiguous between these two epistemic relations, or kinds of knowledge.

While there may be something intuitively compelling about this strategy, we don’t think it accounts for the importance we place in the concept of knowledge. Unless knowledge is knowledge₁, it won’t have the central epistemological position it should have. Knowledge₁, unlike knowledge₂, gives us a guarantee of the importance of the concept involved. It sets a meaningful lower bound on strength of epistemic position: your epistemic position regarding p must be strong enough to make it rational for you to act as if p is true. Therefore, whatever the more important epistemic concept is, purism is false of it.

Appendix

In Section 2, we considered a response to KB that makes it out to be analogous to a conversational implicature. Our main argument against this response was that KB-implications are not cancelable, and so are not conversational implicatures. Here, we provide a more detailed argument that pragmatic construals – whether conversational or otherwise – of KB-implications are implausible. We begin with a worry about whether all conversational implicatures are cancelable.

Grice (1967: 37) distinguishes “particularized” from “generalized” conversational implicatures. The former depend on context in a way that the latter do not (Grice 1967: 37). If I am looking for my piece of pie, and you say to me, “Little Charlie is smiling,” you implicate that Charlie ate the pie. But calculation of this inference depends essentially on awareness of the features of the context of utterance. In other contexts (say if we are posing for a picture), it clearly does not have the same implication. By contrast, consider “I lost a contact lens.” This implicates, in almost all contexts, that the speaker lost his own contact lens.²⁷ No calculation is necessary, either on the part of the speaker in generating the implicature or the hearer in recognizing it.²⁸ The distinction between particularized and generalized implicatures bears on our discussion, because, if KB implications are conversational implicatures at all, then they are generalized ones. Special features of context seem to be irrelevant to the generation and recognition of the implication *S is rational to do A* from saying “S knows that A is the best thing she can do.”

Grice himself (1989: 44) thought that generalized conversational implicatures were cancelable, but this claim has recently been challenged. Patrick Rysiew (2001) argues that

²⁷ Though not, of course, in all contexts. Consider an employee in a contact lens factory doing his end of the day inventory.

²⁸ Levinson (1983, 118-66) discusses a number of possible revisions needed to Grice’s theory of conversational implicatures, many bearing on generalized implicatures. Some of what Grice classified as generalized implicatures may be better classified as “default implicatures.” See note 31.

uncancelability, or at least “uncomfortable” cancelability, can in some cases be explained by an implication’s being nearly universal. The attempted cancelation, “They fell in love and got married but not in that order,” is uncomfortable, he claims, because “and” is used nearly universally to imply temporal ordering (495-6).

We find this criticism of the cancelation test unconvincing. Near universality of an implication, by itself, tells us only about the types of situations in which we usually find ourselves. But it is hard to see why we should be unable, in virtue of our knowledge of a sentence’s meaning, to see how it may be used in situations that are unusual but still not in any way bizarre. As mentioned above, we almost always use “I lost a contact lens” to imply that we lost one of our own contact lenses.²⁹ But it is no entailment. It is easily canceled: “I lost a contact lens, but it was my wife’s, not mine.” The same goes for “They fell in love and got married but not in that order.” Rysiew suggests this cancelation may be precious; but what counts is whether it is absurd.³⁰

However, let us grant for the sake of argument that Rysiew is correct in his claim that near-universal implicatures may be uncancelable. So, a pragmatics theorist might invoke the near universality of KB-implications to explain why 1, 2, and 3, as well as 4, 5, and 6, are so peculiar. But how, then, are we to tell whether KB-implications are generalized implicatures? Rysiew’s (497) discussion suggests a precedence test.³¹ If “p but not q” and “not q but p” are both uncomfortable, then this is evidence that the implication *that q* of saying “p” is not a generalized implicature.³² For example, “John is married but is a bachelor” is every bit as uncomfortable as “John is a bachelor but is married.” Unfortunately for our objector, the

²⁹ Thanks to Anthony Everett for this example and for guidance on the material in this paragraph.

³⁰ Thus, see Levinson (1983 108, 119).

³¹ See also Levinson (1983 98).

³² Rysiew thinks of this precedence test as a negative test for semantic status. If changing precedence removes the implication, it isn’t semantic.

precedence test provides no traction against KB, as is clear from an examination of the pair

(1)/(1R):

(1) Sue knows that going to work is the best thing she can do, but she is not rational to go to work.

(1R) Sue is not rational to go to work, but she knows that going to work is the best thing she can do.

KB-implications thus seem not to be conversational implicatures or any other sort of cancelable implication.³³ There are, however, pragmatic implications that are uncancelable, regardless of precedence. When we assert that *p*, for example, we pragmatically imply that we believe that *p*. And, at least for some properties in some contexts, when we assert that a person has the property, we pragmatically imply that we have the authority to bestow the property on the person (e.g. “You are free to go.”). The uncancelability of these implications, regardless of precedence, is clear from Moore’s paradox (“*p*, but I don’t believe that *p*”) and a similar “performative paradox” (“You are free to go, but I don’t have the authority to release you”).

Might KB-implications fall into this category?

In the two examples given, in performing a speech act, one implies that the performance succeeds, and therefore that it satisfies all the conditions of success, including conditions that are not entailed by what is said. When you assert that *p*, you present yourself as sincere, and so as at least believing that *p*. Similarly, when a police officer tells you that you are now free to leave,

³³ Based on recent work in linguistics (Levinson (1983) and (2000) and Landman (2000) among others) François Recanati (2003 304) distinguishes generalized conversational implicatures from “default” implicatures. Default implicatures are unlike conversational implicatures in a number of ways: they are detachable (by using different words to say the same thing, one can say the same thing without generating the implicature); they arise automatically (rather than through Gricean computation from what is said); and they are embeddable (they come under the scope of logical and other operators). However, default implicatures are similar to conversational implicatures in being cancelable (hence the term “default”). Recanati (2003 305) explains that what begins life as a particularized conversational implicature may become generalized and finally become fully conventionalized into a default implicature. The use of the scalar ‘three’, for instance, is associated with an implication of *exactly three*: “Either John has three kids or he has four, I forget which.” Yet this implicature is cancelable. There is nothing untoward about saying, in response to some who says “John has three kids,” “Well, he does have three kids; and in fact he has four.”

KB-implications, because uncancelable, are not default implications.

she implies that conditions are felicitous for her performative, and so that she has the necessary authority. The conversational maxims underlying these implications are presumably so essential to intelligible conversation that speakers cannot admissibly flout them.

Although first person present-tense knowledge attributions are in many ways special, KB applies in third-person cases as well. It seems far-fetched to say that I perform a special sort of speech act in making third-person knowledge-ascriptions. More importantly, however, there is a further test that distinguishes KB from Moorean and related kinds of what we will call “speech act” implications. Embed the conditional corresponding to the given implication under a necessity operator and ask if it sounds true. Contrast 10 and 11 with 12:

(10) Necessarily, if p, then I believe that p.

(11) Necessarily, if you are free to go, then I have the authority to release you.

(12) Necessarily, if you know that A is the best thing you can do, then you are rational to do A.

10 and 11 sound clearly false, while 12 sounds true. This is evidence that KB-implications are not speech act implications. If they were, then since the content of a speech act can come apart modally from the conditions of one’s successfully making the speech act, we should expect 11 to sound false.

There is a further kind of uncancelable pragmatic implication. *Conventional* implicatures are non-negotiably built into the meanings of certain terms.³⁴ Perhaps it is built into the meaning of ‘but’ that any correct use of “p but q” implicates some sort of contrast between p and q. Other purported conventional implicature indicators include ‘even’ and ‘therefore.’ Asserting “Even A

³⁴ The category of conventional implicature, unlike that of conversational implicature, has never enjoyed widespread acceptance among philosophers of language. Even Grice, who introduced the category, expressed misgivings about it. The main difficulty is that it is hard to see how implications built into the meanings of words, and come under the scope of logical operators, could fail to be semantic. See Bach (1999) and Grice (1989). We should add, however, that the category of conventional implicature seems to be staging a comeback recently. See Barker (2003) and Potts (2003).

is F” implicates that there are Fs other than A and that it is surprising or remarkable that A is F. Asserting “p, therefore q” implicates that q is inferable from p.³⁵ Perhaps ‘knows,’ too, has associated conventional implicatures. The pragmatics theorist might claim that it is built into its meaning that any assertion of ‘S knows that p’ implicates that S is rational to act as if p.³⁶

One might think that modal embedding considerations would clearly distinguish KB-implications – and perhaps KA-implications as well – from conventional implicatures. Compare 12 above with 13 and 14:

(13) Necessarily, if someone is small but powerful, then there is some contrast between being small and powerful.

(14) Necessarily, if even the Rockies have flat spots, then it is surprising that the Rockies have flat spots.

While 12 sounds clearly true, neither 13 nor 14 do. But nor do they sound clearly false. They simply sound bizarre.³⁷ This might be some evidence that KB-implications are not conventional implicatures, but we do not wish to rest too much on it. There is a further and better test that will enable us to distinguish KB- and KA-implications from conventional implicatures.³⁸

³⁵ Christopher Potts (2003) has recently argued that supplements (e.g., appositives, *As*-parentheticals) and expressives (e.g., epithets) are associated with conventional implicatures. E.g., saying “Quine, a brilliant philosopher, attended Oberlin” implies that Quine was a brilliant philosopher, even though what is said does not entail this implication. Potts expresses skepticism about whether there are words that carry both truth-conditional content and implicatures. Thus, he denies that conventional implicatures are associated with the standard examples of ‘but’, ‘even’, etc.

³⁶ If “S knows that p” conventionally implicates *S is rational to act as if p*, then its implicature behaves differently than do those of ‘but’, ‘even’, and ‘therefore.’ When ‘knows that p’ is embedded in an unasserted context, the implication of rationality to act is *never* projected to the whole utterance. In this way, the relation between ‘knows that p’ and ‘is rational to act as if p’ is exactly the same as the relation between expressions and their entailments.

³⁷ Perhaps the source of the oddity in 12 and 13 is that the purported implicature content associated with words like ‘but’ and ‘even’ can come under the scope of logical operators. Conventional implicatures are not always projected from an embedded material to the utterance as a whole. For example, as Barker notes (2003 9), in saying, “It is not that she dislikes him *but* dates him *nevertheless*. She dates him because she dislikes him,” one does not imply that there is a contrast between her disliking him and her dating him.

Barker’s examples can be disputed. For example, one might think that the negation involved in the example cited is “metalinguistic” negation, commenting on the inappropriateness of saying something or of saying it in a certain way. We will not enter into this debate. Suffice it to say that we will not rest our argument on the differential oddity between 11 and 12/13.

³⁸ One might think that there a quick argument that KA- and KB-implications aren’t conventional implicatures. Conventional implicatures are detachable, but KA- and KB-implications aren’t. One can replace “knows that” with

We will call this the “semantic ascent” test (Grice 1967 25-6, Barker 2003 14). Try to think of situations in which *q* clearly does not obtain, but about which you would be willing to say, if asked, that “*p*” is true in that situation.³⁹ If after suitable reflection you cannot think of any, this is evidence that the implication is not a conventional implicature. If, in addition, you find yourself, upon considering an arbitrary not-*q* situation, feeling strong pressure to call “*p*” false, this is even stronger evidence that the implication is not a conventional implicature. The rationale for this “semantic ascent” test is clear. Conventional implicatures are not part of truth-conditional content, nor are they entailed by truth-conditional content. So, it should be possible for the truth-conditional content to obtain without the implicature content. Semantic ascent avoids the problems that we saw with embeddings because calling a statement true in an actual or hypothetical situation does not commit one to the conventional implicatures generated by asserting the statement in an actual or hypothetical situation. Because pragmatic implications are not entailments, the semantic ascent test in fact serves as a negative test for pragmatic implication in general, and not only for conventional implicature.⁴⁰

“is aware that,” “realizes that,” “recognizes that,” “sees that,” etc. and the implication remains. Unfortunately, detachability tests are often inconclusive, when semantic content is in dispute (see Bach 1999), as illustrated in the following dispute about ‘but’:

A: The contrast implication associated with ‘but’ is non-detachable. Whether one uses ‘despite’, ‘in contrast to’, ‘although’, etc. the implication remains.

B: ‘And’ will not preserve the implication.

A: That’s question-begging. I think ‘and’ and ‘but’ have different semantic contents.

³⁹ Similar tests are employed in Grice (1967), DeRose (2002), and Barker (2003). (DeRose’s “check the negations” test, however, is not metalinguistic.) One might wish to require that the relevant situations not be too remote from actuality (e.g., by building in substantially different laws of nature).

⁴⁰ Here are some examples of the application of the semantic ascent test. The test rightly gives the result that the implication *S is a man* generated by saying, “*S is a bachelor*” is not pragmatic. Describe any situation in which *S* is not a man, and none of us will say that “*S is a bachelor*” is true in it. Consider next whether “Even Carville likes Clinton” is true, as things actually are. The sentence sounds strange, but we are presumably willing to call it true if pressed. It certainly contrasts with the likes of “Even Ken Starr likes Clinton,” which is false. The same goes for the purported conventional implicatures associated with ‘but’, as well as with paradigm supplements (“W.V.O. Quine, the great defender of the a priori, attended Oberlin” – true, but odd) and expressives (“The evil Yankees won again” – again, true, but odd, at least for Yankees fans). There are hard cases, however. Is “Even George Bush is married to Laura Bush” true? It’s difficult to say, but we do not feel strong pressure to call it false (as we would if, say, the Bushes divorced). So, the semantic ascent test, at least, does not provide strong evidence against the hypothesis that the implication *A is not the only F* from “Even *A is F*” is pragmatic.

What does the semantic ascent test have to say about KB- and KA-implications? We briefly noted in Section II.1 that we reason in conformity with the contrapositive of KA. Here we make a similar point. When we deny that a subject is rational to act as if *p*, we feel considerable pressure to deny that he knows that *p*. Or, in terms of KB, when we deny that a subject is rational to do *A*, we feel considerable pressure to deny that the subject knows that *A* is the best thing she can do. In fact, we feel pressure to call the claim “*S* knows that *A* is the best thing she can do” false. About Sue in the *a/c* case, for example, we will not be willing to say, “She knows that going to work is the best thing she can do” is true; we will in fact feel strong pressure to call it false. Therefore, the semantic ascent test provides strong evidence that KB-implications – and so also KA-implications – are not pragmatic.

We have seen that KB-implications are not conversational implicatures (particularized or generalized), because they are uncancelable, regardless of precedence. We have also seen that they are not species of the two best-known kinds of uncancelable pragmatic implication: speech act implications and conventional implicatures. Unlike the former, KB-implications seem to yield truths when embedded under a necessity operator. Unlike the latter, and, in fact, unlike pragmatic implications in general, KB-implications are deemed non-pragmatic by the semantic ascent test. Our discussion thus provides strong evidence that KB-implications are not pragmatic, and so, by extension, that neither are KA-implications. The pragmatics theorist, thus, faces an uphill battle. She must tell us what sort of pragmatic implication KB and KA could involve if not one of the familiar kinds we have discussed, and she must provide evidence that this unfamiliar kind of implication is indeed pragmatic. Unless and until she can discharge these responsibilities, we see no reason not to take KB to be valid and KA to be true.

Bibliography

- Bach, Kent. 1999. "The Myth of Conventional Implicature." *Linguistics and Philosophy*, 22:327-66.
- Barker, Stephen. 2003. "Truth and Conventional Implicature." *Mind* 112, 445:1-33.
- Cohen, Stewart. forthcoming. "Knowledge, Speaker, and Context." *Philosophical Issues*.
- _____. 1999. "Contextualism, Skepticism, and the Structure of Reasons." *Philosophical Perspectives* 13:57-89.
- _____. 1998. "Contextualist Solutions to Epistemological Problems: Skepticism, Gettier, and the Lottery." *Australasian Journal of Philosophy*, 76, 289-306, and reprinted in Sosa and Kim's anthology *Epistemology*, 517-530, Oxford: Blackwell, 2000.
- DeRose, Keith. 2004. "The Problem with Subject-Sensitive Invariantism." *Philosophy and Phenomenological Research* 68:346-350.
- _____. manuscript. "On the Ordinary Language Basis for Contextualism and the New Invariantism."
- _____. 2002. "Assertion, Knowledge, and Context." *Philosophical Review* 111, 2:167-203.
- Fantl, Jeremy and McGrath, Matthew. 2002. "Evidence, Pragmatics, and Justification." *The Philosophical Review* 111, 1:67-94.
- Grice, Paul. 1967. "Logic and Conversation." William James Lectures at Harvard. Reprinted in Paul Grice *Studies in the Ways of Words*, 1989, Cambridge, MA: Harvard University Press: 22-40.
- _____. 1989. "Further Notes on Logic and Conversation." In Grice 1989, 41-57.
- Hawthorne, John. 2004. *Knowledge and Lotteries*. New York: Oxford.
- Hofweber, Thomas. 1999. "Contextualism and the Meaning-Intention Problem" in *Cognition, Agency, and Rationality*, ed. by Korta, Sosa, Arrozola.
- Hookway, Christopher. 1990. *Scepticism*. London: Routledge.
- Landman, Fred. 2000. *Events and Plurality* (Dordrecht: Kluwer).
- Levinson, Stephen 2000. *Presumptive Meanings: The Theory of Generalized Conversational Implicature* (Cambridge, Mass: MIT Press).
- _____. 1983. *Pragmatics* (Cambridge: Cambridge University Press).
- Lewis, David. 1996. "Elusive Knowledge." *Australasian Journal of Philosophy* 74: 549-67.
- Nozick, Robert. 1993. *The Nature of Rationality*. Princeton: Princeton University Press.

- Owens, David. 2000. *Reason without Freedom: The problem of epistemic normativity*. London: Routledge.
- Potts, Christopher. 2003. *The Logic of Conventional Implicatures*. Dissertation. University of California – Santa Cruz.
- Recanati, François. 2003. “Embedded Implicatures.” *Philosophical Perspectives, 17, Language and Philosophical Linguistics*: 299-332.
- Rysiew, Patrick. 2001. “The Context-Sensitivity of Knowledge Attributions.” *Nous* 35:477-514.
- Stanley, Jason (forthcoming). *Knowledge and Interest*. Oxford University Press, Lines of Thought Series.
- Williams, Michael. 1992. *Unnatural Doubts*. Oxford: Blackwell.