

Fallibilism, epistemic possibility, and concessive knowledge attributions

Trent Dougherty, University of Rochester

Patrick Rysiew, University of Victoria

I. If knowing requires believing on the basis of evidence that entails what's believed, we have hardly any knowledge at all. Hence the near-universal acceptance of fallibilism in epistemology: if it's true that "we are all fallibilists now" (Siegel 1997: 164), that's because denying that one can know on the basis of non-entailing evidence¹ is, it seems, not an option if we're to preserve the very strong appearance that we do know many things (Cohen 1988: 91). Hence the significance of *concessive knowledge attributions* (CKAs) (Rysiew 2001) – i.e., sentences of the form 'S knows that *p*, but it's possible that *q*' (where *q* entails not-*p*). To many, utterances of such sentences sound very odd indeed. According to David Lewis (1996: 550), however, such sentences are merely "overt, explicit" statements of fallibilism; if so, their seeming incoherence suggests that, contrary to our everyday epistemic pretensions, "knowledge must be by definition infallible" after all (*ibid.*: 549).

Recently Jason Stanley (2005) has defended fallibilism against the Lewisian worry that overtly fallibilistic speech is incoherent. According to Stanley, CKAs are not just odd-sounding: in most cases, they are simply false. But this doesn't impugn fallibilism. Insofar as the odd-sounding utterances Lewis cites state the fallibilist idea, the latter portion thereof ('S cannot eliminate a certain possibility in which not-*p*', e.g.) expresses the idea that the subject's evidence doesn't entail what's (allegedly) known (hence, the negation of any contrary propositions).

¹ Stanley (2005: 127) defines fallibilism as "the doctrine that someone can know that *p*, even though their evidence for *p* is logically consistent with the truth of not-*p*"; Cohen (*op cit.*) defines it as the view that one can know on the basis of non-entailing reasons. Here, we'll use the above way of stating the same idea. For a good concise discussion of fallibilism, see Feldman 1981.

According to Stanley, however, this is not the best reading of the possibility clauses CKAs contain. On the correct account of the latter, while the sentences Lewis cites are almost always self-contradictory, they don't capture the fallibilist idea after all. Here, we argue that the sentences in question do express precisely the fallibilist idea, but argue that Lewis has nonetheless failed to raise a problem for the latter. In addition, we respond to worries that the resulting view of the semantics of epistemic possibility statements has certain unacceptable consequences.

2. Odd-sounding CKAs can take many forms:

- (1) (a) I know that Harry is a zebra, but it's possible that Harry is a painted mule.
- (b) John knows that Harry is a zebra, but it's possible that Harry is a painted mule.
- (c) John knows that Harry is a zebra, but it's possible for John that Harry is a painted mule. (Stanley 2005: 126.)

Lewis' own attempt "to thread a course between the rock of fallibilism and the whirlpool of scepticism" (Lewis 1996: 566) involves embracing epistemic contextualism: we should regard "knows" either as expressing different relations in different contexts or as a tacit indexical. Having done so, we may say with the infallibilist that *S* knows that *p* iff *S*'s evidence eliminates every possibility in which not-*p* (*ibid.*: 551). But since "every" is (Lewis claims) restricted to a particular conversational domain, and since certain not-*p* possibilities will be "properly ignored" in a given situation (*ibid.*: 553-4), we preserve our intuitive non-scepticism.

For many, though, accepting contextualism is itself a pretty hefty price to pay: it would be much better to resolve the problem Lewis has set up by some other means. As Stanley says, then, "the challenge is to explain why it seems that *speaking* of fallible

knowledge is odd, without impugning the truth of fallibilism, or adopting contextualism” (2005: 126).

A previous attempt to meet this challenge is Rysiew’s (2001: 492-4). Rysiew claims that an utterance of ‘It’s possible that p ’ typically pragmatically conveys that the speaker doesn’t know that $\text{not-}p$, and doesn’t know that p . In the case of CKAs, however, the speaker explicitly contradicts the latter information: hence their oddity. Since the explanation is purely pragmatic, Rysiew allows that CKAs themselves may often be true. Stanley finds this result unacceptable:

...Rysiew claims that the fact that uttering any of the sentences in (1) results in oddity is due to *just* pragmatic facts, namely an implicature associated with epistemic possibility statements. This entails that he believes that it can in fact be epistemically possible for a speaker that q , even though that speaker knows that $\text{not-}q$. This is an unacceptable consequence. The oddity of most utterances of these sentences has something to do with the *semantics* of epistemic possibility statements. (2005: 127.)

According to Stanley, that semantics is best rendered as follows (*ibid.*: 128):

(EPk) It is possible for A that p is true if and only if what A knows does not, in a manner that is obvious to A , entail $\text{not-}p$.

Equipped with (EPk), we can retranslate the problematic CKAs:

- (2) (a) I know that Harry is a zebra, but what I know does not entail, in a manner that is obvious to me, that Harry is not a painted mule.
- (b) John knows that Harry is a zebra, but what I know does not entail, in a manner that is obvious to me, that Harry is not a painted mule.

(c) John knows that Harry is a zebra, but what John knows does not entail, in a manner that is obvious to him, that Harry is not a painted mule. (*Ibid.*)

These are pretty clearly false in most² instances – at least, so long as the subject (the speaker, or John) is “clearly minimally intelligent” (*ibid.*).³ But, as we can now see, all the work of explaining the alleged infelicity of (1a)-(1c) can be done without raising even a *prima facie* problem for fallibilism; for, whatever the source of the oddity to which utterances of (2a)-(2c) might give rise, since the latter half of these claims concern what is/isn’t entailed by what one knows, rather than what is/isn’t entailed by one’s evidence, none of those sentences express the fallibilist idea.

3. Assume that our conversations are governed by the mutual presumption that others try to conform to Grice’s (1989) Co-operative Principle (CP) – hence, that they strive to make their conversational contributions maximally relevantly informative.⁴ Then, if what’s at issue is whether p , if the speaker takes the subject (himself, or another) to know either that p or that not- p , that’s would he should say, since either entails a univocal answer to the question at issue; whereas, on anyone’s view, ‘It’s possible that p ’ does not. So, from a speaker’s saying the latter, one may infer that he doesn’t take himself to know either that p or that not- p .

Given that CKAs are odd-sounding, in the absence of a plausible deflationary explanation of that fact, it might be reasonable to regard them as semantically confused. But, as even the foregoing preliminary argument shows, given only quite minimal assumptions about their

² Stanley has reservations about saying that (1b) and (2b) are false, and prefers a pragmatic account of their infelicity (2005: 128ff.) Plausibly, though, (2b) – and thus, on Stanley’s account, (1b) as well -- will be false in most instances, since in most instances if one has any justification for thinking that John knows that Harry is a zebra, it will suffice for one’s knowing the latter oneself, and the latter obviously entails that Harry isn’t a painted mule.

³ Intelligent enough to see, e.g., that something’s being a zebra obviously entails that it’s not a painted mule (*ibid.*).

⁴ This is Mike Harnish’s (1976) proposed compression of the various Maxims and sub-Maxims falling under the CP.

semantics we have every reason to think that CKAs are *at least* pragmatically odd: the speaker's not knowing that p is inferable from his saying "It's possible that not- p ," but that he does know this is what the first half of CKAs explicitly affirm. On its own, then, the odd-sounding character of CKAs does not require EPk; and so long as EPk is not presumed, it's an open question whether the correct semantics of epistemic possibility statements must have it, as Stanley assumes, that 'It's possible (for S) that not- p ' entails ' S doesn't know that p ' – hence, that CKAs are not just pragmatically odd, but semantically defective as well. Aside from its conflicting with the latter assumption, Stanley's reason for rejecting Rysiew's purely pragmatic explanation of CKAs' oddity is simply "that it is mysterious to me what he takes the *semantic content* of epistemic possibility statements to be" (2005: 127). As Stanley says, "[t]he correct characterization of epistemic possibility is a vexed matter" (*ibid.*). One natural and obvious suggestion, however, is that required by Lewis' argument: the concession clauses in CKAs express precisely the fallibilist idea; what is epistemically possible for a subject are those things which his *evidence*, rather than what he *knows*, does not rule out. Thus,

(EPev) (a) q is epistemically possible for S iff not- q isn't entailed by S 's evidence;

or, equivalently, that

(b) q is epistemically possible for S iff q has non-negligible probability on S 's total evidence.^{5,6}

Assuming EPev, (1a)-(1c) become:

⁵ Just as EPk *per se* is neutral about what knowledge is, EPev itself is neutral as to what counts as evidence – whether it includes, or is restricted to, one's experiences and (apparent) memories, one's beliefs, and so on. For the same reason, EPev doesn't require a specifically internalist handling of evidence. Maintaining the distinction between EPk and EPev does require rejecting Williamson's (2000) principle that 'E=K' – i.e., that one's evidence is one's knowledge. But that principle is sufficiently controversial that most would join us in assuming that the distinction between EPk and EPev is real.

⁶ Something like this appears to be Lewis' view of epistemic possibility, which he couples with a particular view of evidence: " S 's epistemic possibilities are just those possibilities that are uneliminated by S 's evidence....[T]he uneliminated possibilities are those in which the subject's entire perceptual experience and memory are just as they actually are" (1996: 552-3). For similar ideas, see Kripke 1972: 103-4 and 142-3.

- (3) (a) I know that Harry is a zebra, but there is a non-negligible probability on my total evidence that Harry is a painted mule.
- (b) John knows that Harry is a zebra, but there is a non-negligible probability for me that Harry is a painted mule.
- (c) John knows that Harry is a zebra, but there is a non-negligible probability for John that Harry is a painted mule.

Since, given fallibilism, knowledge does not require evidence that entails what one knows, and since our beliefs are in fact rarely based on evidence that entails what's believed, EPev has the result that sentences like (3a)-(3c) are not only consistent, but will often express truths. This is a point in favor of EPev, as against EPk. As we've seen, the latter approach has the result that such sentences are very often false, indeed contradictory. This leaves unexplained why, apart from their simply being confused, competent speakers would ever utter such sentences, which they clearly do (a point to which we'll return below). On the current view, however, when people utter CKAs they are rightly expressing the fact that, strictly speaking, their evidence leaves open the possibility of error: they are simply acknowledging the sober truths of *fallibilism* -- that knowledge is compatible with evidence that doesn't entail what's believed -- and *fallibility* -- that we very rarely have evidence which entails our conclusions.

Of course, these same truths together with EPev entail that *very* many things will be epistemically possible for a subject. But the obvious truth in most instances of 'It's possible that *p*' does not make it not worth saying. On the contrary, just because something's being epistemically possible will in most cases be mutually obvious, saying so will be *prima facie* infelicitous. This may be used to bolster the preliminary argument given above, as to the source of any discomfort attending CKAs.

Thus, even though, assuming EPev, ‘It’s possible [for *S*] that not-*p*’ does not imply ‘*S* doesn’t know that *p*’, the latter is typically inferable from one’s saying the former: given the CP, one’s saying that not-*p* is possible implies, not just (what EPev states) that not-*p* has non-negligible probability, but that one has some real grounds for supposing not-*p* might be the case and that one (therefore) isn’t confident that *p*. But if that’s so, one shouldn’t claim to know that *p*, or attribute knowledge to another, since in so doing one of course represents oneself as (confidently) believing that *p* and as having adequate evidence for that belief. And if one *doesn’t* think there’s not merely a non-zero chance that not-*p* – that not-*p* is not *merely* possible – why mention it? Under the circumstances, doing so would be misleading at best. In more general terms, either the doubt or reservation which “it’s possible that not-*p*” is naturally understood as indicating is significant, or it is not. If it is, there’s a norm to hedge the assertion which comprises the first half of CKAs. This may be a generic consequence of the CP, or a consequence of the Maxim of Quality.⁷ If, however, the doubt is not significant, then the Maxim of Relation recommends that one not mention it. Either way, the explanation of the oddity of CKAs is pragmatic.

4. John Hawthorne also subscribes to something very like EPk.⁸ He says “It sounds extremely odd to assert ‘I know that *p* but there is a chance that not-*p*’ or ‘I know that *p* but there is a chance that *q*’ (where *q* is known to be incompatible with *p*)” (Hawthorne 2004: 24). However, to the contented fallibilist, any strangeness there might have been dissipates once we have a good handle on the considerations just rehearsed. In fact, Hawthorne himself notes (*ibid.*: n. 60) that

⁷ ‘Do not say what you believe to be false, or for which you lack adequate evidence’ (Grice 1989: 27).

⁸ “It is possible that *p* for *S* at *t* (There is a chance that *p* for *S* at *t*) iff *p* is consistent with what *S* knows at *t*” (2004: 26).

others⁹ have pointed out to him exceptions to the generalization he makes -- “I know they are going to lose but I’m going to carry on watching just in case,” and “I know it’s not going to rain, but still, I’m going to take an umbrella just in case.” Hawthorne claims, “Such utterances do have an air of retraction about them, however.” But that’s contentious as well. For instance, the disutility of getting soaked multiplied by the small probability of its raining could easily have an absolute value greater than the utility of not having to put the umbrella in your bag. These are commonplace considerations.

However, if there is a natural awkwardness, it needs to be explained. Hawthorne considers the possibility of a pragmatic explanation along the lines of the one presented above: “Perhaps the reason why ‘I know that p , but there is a chance that not- p ’ is unassertable is that it is already obvious to audiences (for nearly any p) that there is always a chance that not- p ” (*ibid.*: 25). But mutual obviousness doesn’t make for unassertability: there are plenty of patently true sentences -- “Tomorrow’s another day,” “People do sometimes lie,” “You can’t change the past,” etc. -- that we regularly utter to good illocutionary (and perlocutionary) effect. Nevertheless, the seeming strangeness of certain CKAs would need explaining. But, as we argued above, the (typical) mutually obvious banality of possibility statements forms part of a pragmatic explanation of the phenomenon -- not to mention, of how we are able to put other mundane sentences to good communicative use.

Hawthorne replies: “However, this doesn’t happily accommodate the fact that we can respond to such an assertion, not by complaining that too much has been said, but by noting that the claim is self-defeating: ‘Well, if there is a chance that not- p , then you don’t know that p , do you?’” (*ibid.*: 25). Our noting this, however, is easily explained on the present view. In saying that p is possible one implies that there is some real chance that p . Given that one’s audience

⁹ Tamar Gendler and Brian Weatherson.

realizes this, it will be only natural for them to respond in kind – i.e., by referring to a *significant* chance that p as, simply, “a chance” that p . And a significant chance of error may well be enough to prevent one from knowing.

Hawthorne adds this to his rebuttal: “[T]he approach would predict that the speech ‘I know that p . There’s no chance that not- p ’ would be unassertable for a different reason -- by containing an obvious falsehood. But such speeches are commonplace” (*ibid.*). But here too, there’s no reason for the proponent of EPev to predict the unassertability of such sentences. Just as a sentence’s banality needn’t prevent it from being used to good effect, neither must its literal falsity: just as, in saying that p is possible one implies that there is some real chance that p , one’s saying that there’s no chance that p may be understood as a piece of communicatively effective hyperbole. For again, given the maxim of Relation, in saying “there’s no chance that p ” one is naturally taken to mean that there’s no *significant* chance that p .¹⁰ As stated above, that speakers are regarded as uttering falsehoods is something generally to be avoided where we’re given no reason why they would do so. However, in the present case, such a reason has been provided.

It bears emphasizing that, contrary to what one might expect given the EPk-theorist’s diagnosis of the oddity of CKAs, we *do* utter such sentences – often, without any attendant oddity. Of course, we rarely do so in the straight and abstract form upon which philosophers tend to focus – the simple “I know that p , but it’s possible that not- p ,” for example.¹¹ Rather, when we express the information embodied in CKAs, we tend to say things like, “Of course it’s *possible* that oil prices will fall dramatically over the next month, but we all know that that’s not going to happen” (cf. Rysiew 2001: 497). If EPk is correct -- if it’s true that, if there’s an epistemic possibility that not- p , then you don’t know that p -- then less direct and more naturally phrased

¹⁰ Hawthorne (2004: 25, n. 63) anticipates this move, but does not develop or endorse it.

¹¹ Though we do sometimes do even that: a referee for this paper recommended our making a certain small point explicit, saying, “I know no one will get confused about this, but it’s possible someone will.”

CKAs such as this should also come out as false, and as easily recognizable as such, for any “minimally intelligent” person. But this is not the natural result. Or consider this example: “We now know that there is a top quark: we have presented considerable evidence indicating its existence. Of course, our experiments are not conclusive, and there is a chance that our results will be overturned. We are confident, however, that they will be borne out.” This is a natural way for a scientific article to end. The authors are clearly asserting that there is a top quark and that they know this; but they also explicitly allow the possibility that they are wrong. This is not irrationality, or a blatant flouting of an obviously correct view of epistemic possibility. What they say is just good common sense. Fallibilism underwrites its truth, and EPev provides the semantics for it.

To sum up, like Stanley’s response to Lewis based on EPk, the present response based on EPev shows that the non-contextualist fallibilist has nothing to fear from the data Lewis cites. Further, unlike Stanley’s account, EPev does not involve imputing to actual utterers of CKAs any unexplained gross failures of intelligence or linguistic competence. Also, we have avoided requiring at the outset that the problem set by Lewis must have a specifically semantic solution,¹² and have responded to Hawthorne’s objections to the sort of view we offer. Lastly, in offering a pragmatic resolution of the problem Lewis sets, we have undercut one of the principal motivations for EPk itself – viz., the apparent inconsistency of, e.g., ‘It’s possible that not-*p*’ and ‘I know that *p*’, and the absence of any plausible deflationary (i.e., non-semantic) explanation of such a clash.¹³

¹² Cf. ‘Grice’s Razor’ (Davis 2005): “Other things equal, it is preferable to postulate conversational implicatures rather than senses, conventional semantic implicatures, or semantic presuppositions because conversational implicatures can be derived from independently motivated psycho-social principles.”

¹³ See, e.g., DeRose 1991: 596-601 and Hawthorne 2004: 24-28.

References

Cohen, S. 1988. How to be a fallibilist. *Philosophical Perspectives*, 2, *Epistemology*: 91-123.

Davis, W. Implicature. *Stanford Encyclopedia of Philosophy*, E. N. Zalta, ed.:

<http://plato.stanford.edu/entries/implicature/>.

DeRose, K. 1991. Epistemic possibility. *The Philosophical Review* 100.4: 581-605.

Feldman, R. 1981. Fallibilism and knowing that one knows. *The Philosophical Review* 90.2: 266-282.

Grice, H. P. 1989. *Studies in the Way of Words*. Harvard University Press.

Harnish, R. 1976. Logical form and implicature, *An Integrated Theory of Linguistic Abilities*, T. Bever, J. Katz, and T. Langendoen, eds. Crowell.

Hawthorne, J. 2004. *Knowledge and Lotteries*. Oxford University Press.

Kripke, S. 1972. *Naming and Necessity*. Harvard University Press.

Lewis, D. 1996. Elusive knowledge. *Australasian Journal of Philosophy* 74: 549–567.

Rysiew, P. 2001. The context-sensitivity of knowledge attributions. *Noûs* 35.4: 477–514.

Siegel, H. 1997. *Rationality Redeemed?* Routledge.

Stanley, J. 2005. Fallibilism and concessive knowledge attributions. *Analysis* 65.2: 126–131.

Williamson, T. 2000. *Knowledge and Its Limits*. Oxford University Press.