

## Clarity about Concessive Knowledge Attributions: Reply to Dodd<sup>i</sup>

Trent Dougherty

Patrick Rysiew

Recently, Dylan Dodd<sup>ii</sup> has tried to clear up what he takes to be some of the many confusions surrounding concessive knowledge attributions (CKAs) -- i.e., utterances of the form “S knows that  $p$ , but it’s possible that  $q$ ” (where  $q$  entails not- $p$ ) (Rysiew 2001). We agree with Dodd that there are a lot of mistaken beliefs amongst philosophers who’ve discussed CKAs. However, we fear that in his criticisms of Dougherty and Rysiew (2009), Dodd has created more clutter than clarity himself. In this brief note, we tidy up a bit with concise replies to the three points of criticism (and related matters) that Dodd raises.

### *I. Grice Misunderstood?*

Dodd’s first charge is that we misunderstand some elementary facts of Gricean pragmatics. The charge concerns our explanation of the sometime<sup>iii</sup> infelicity of CKAs. We claimed that the typical oddity of CKAs is hardly surprising: one reason utterances of the form “I know  $p$ , but it’s possible that not- $p$ ” frequently sound odd is that utterances of ‘It’s possible that  $p$ ’ generally pragmatically convey<sup>iv</sup> *I don’t know whether  $p$  is the case (I don’t know that  $p$  and I don’t know that not- $p$ )*. Thus, whatever the correct semantics for epistemic possibility statements, the second conjunct of the CKA will convey the negation of the first. Naturally this can result in discordance in the ear of one’s interlocutor; so we have every reason to expect that CKAs will often be *at least* pragmatically infelicitous. We then move beyond his preliminary argument and present a semantics for expressions of epistemic possibility – viz., that  $q$  is epistemically possible for  $S$  just in case not- $q$  isn’t entailed by  $S$ ’s evidence or, more finely, just in case  $q$  has non-

negligible/non-zero probability<sup>v</sup> on *S*'s total evidence -- whereby such oddity as attends them is *merely* pragmatic. Briefly: given the facts of fallibilism about knowledge and our own fallibility, we'll often know things the negations of which aren't entailed by our evidence or that have non-zero probability on our evidence;<sup>vi</sup> so CKAs can indeed express truths. Still, while 'It's possible [for *S*] that not-*p*' doesn't imply '*S* doesn't know that *p*', one's saying that not-*p* is possible implies, not just that not-*p* has non-negligible (/non-zero) probability, but that one has some relevant grounds for supposing that 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 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? So goes our account of the merely pragmatic source of the sometime-oddity of CKAs.

But, says Dodd: "What Dougherty and Rysiew didn't seem to notice is that just because an utterance of '*p*' implicates that not-*q* doesn't mean that an utterance of '*p* and *q*' is infelicitous" (p. 389). What Dodd doesn't notice, however, is that we rely on no such assumption as a premise: we do not argue by noting that CKAs are an instance of the schema *utterance of the form 'p and q', where "p" implicates not-q* and then inferring that they must therefore be infelicitous. Some instances of that schema are infelicitous and some aren't. Like most utterances, their felicity varies by context of utterance. Indeed we point this out—as Dodd even notes (Section 6)—as part of our case. But if we think that CKAs are sometimes felicitous, *how could* we be relying on an assumption that instances of the schema are perforce infelicitous? After all, one of the tests of some information's being merely pragmatically generated is cancelability. But an implicature (e.g.; remember n. 4) is canceled precisely by saying '*p* and *q*'

even though ‘*p*’ implicates not-*q*.<sup>vii</sup> Since we argue for the felicity of certain CKAs on the grounds that, in context, they constitute an effective cancelation, it is strange to accuse us of missing this point.<sup>viii</sup>

Dodd illustrates the schema point with the following example:

[S]aying in a letter of recommendation ‘This applicant has wonderful handwriting’ typically implicates that he’s not a very talented philosopher, but that doesn’t mean writing ‘This applicant has wonderful handwriting and is a very talented philosopher’ is infelicitous. (*Ibid.*)

Quite right. But while this might illustrate one’s not being able to infer the infelicity of certain utterances from the schema alone, once again that has no direct bearing upon CKAs or our treatment thereof.

First, notice that the schematic way of representing such cases is misleading in an important way: it’s not the bit about the handwriting *per se* that generates any information about the candidate’s philosophical abilities (or lack thereof); it’s the *silence* on the ostensible subject that does so. Insofar as there is an implicature from “*p*” to not-*q* here, that’s so only *modulo* the context, which includes the assumption that one is supposed to speak to the philosophical talents of *X*. In fact, holding that context fixed, almost *any* assertion could be the value of *p* here -- that, say, *X* has good hygiene or is a good snowboarder. So again, it’s not “*p*” itself but the letter-writer’s silence on the topic of interest that’s doing the real work here. And similarly in the case of those CKAs which are infelicitous, we say: their infelicity isn’t to be explained merely by their instantiating the schema, but by the speaker’s saying what he does given the specific, contextually variable considerations described above. Further, even if it *were* “*p*” all on its own that was generating the implicature in Dodd’s example, all the example would show is that

implicatures may be canceled and that, when they are, the result is felicitous.<sup>ix</sup> But that's not news; and again, it's precisely what we argue is the case with certain suitably chosen CKAs.

## II. "For all I know..."

Second, Dodd says: "there is the following problem with [Dougherty and Rysiew's] proposal. Speakers take 'It's possible that  $p$ ' to be more or less equivalent to 'For all I [or *we*] know,  $p$ '" (*ibid.*). Therefore, he says, we need to explain the infelicity of the following utterances – this is Dodd's "second desideratum", the first being to explain the infelicity of CKAs; and, of course, we need to do so in a way that doesn't assume what Dodd calls "the received view" of epistemic possibility:<sup>x</sup>

- (1) For all I know,  $p$  is true, but  $p$  is definitely false.
- (2) For all I know,  $p$  is true, but  $p$  has to be false.
- (3) For all I know,  $p$  is true, but it's not the case that  $p$  might be true.

Dodd assumes that 'the received view' of epistemic possibility has a significant edge over our own here. Why? Because in (1)-(3) there's mention of knowledge, in which that view trades, so there's a direct route from what the first conjunct expresses to something incompatible with what the second asserts: since, on the received view, epistemic possibility just is (relatively transparent) consistency with one's knowledge, the first half of the phrase would express *possibly*  $p$  and the second half expresses *necessarily not- $p$* . So (1)-(3) are infelicitous because they're straightforwardly self-contradictory. Whereas, there's no mention of evidence (our central notion) in (1)-(3); and neither is anything asserted to be known (p. 393). Dodd infers that whatever our account of the infelicity of (1)-(3), it will be disjoint from our account of the

infelicity of CKAs; so unlike proponents of the received view we can't offer a unified explanation of CKAs and (1)-(3) (pp. 389-390).

But here too, as with the criticism discussed in the previous Section, Dodd has an unduly restrictive view of just what explains the infelicity of CKAs on our account. (In this, he's perhaps abetted by the fact that he presents the preliminary argument for the at-least-pragmatic infelicity of CKAs as though it were our complete account thereof (see Section I, above, and p. 389 of Dodd).) The oddity of CKAs, on our view, is due to utterances thereof conveying inconsistent information. But there's nothing, so to speak, *magical* about "it's possible that not- $p$ " and/or "S knows that  $p$ " such that they, and they alone, might contribute to the generation of the type of mixed message to which many CKAs give rise. Though the conjuncts in CKAs are themselves perfectly consistent, their joint utterance is odd (when it is) because the utterer thereby conveys conflicting information about his credal-epistemic situation – i.e., his degree of confidence in  $p$ , the quality of his evidence as to  $p$ , whether he regards not- $p$  as a serious possibility and so a legitimate grounds for doubt, and so on (see I, above).

Once we're clear about this, it's apparent that our explanation of the oddity of (1)-(3) isn't a different *sort* of explanation at all. In fact, it's not much different from the treatment of (1)-(3) that the received view recommends. Thus, with (1) and (2), the key is that "definitely" and "has to" imply that S's evidence is conclusive (or so regarded by the speaker). And though we ground epistemic possibility in evidential probability rather than knowledge, *conclusive* evidence *entails* being in a position to know (*modulo* non-Gettiered true belief) Thus, someone who says (honestly and soberly) " $p$  is definitely false" represents themselves as having conclusive evidence that it is false, which is to represent themselves as knowing that  $p$  is false. This is clearly in tension, and very likely simply inconsistent, with "For all I know,  $p$  is true."

The third case is really no different. For the second conjunct is equivalent to ‘*p* must be false’ which is the same as in (1) and (2).

In general terms, then, translating (1)-(3) according to our semantics for epistemic possibility, either those sentences themselves or what utterers thereof would represent as being the case all have the following form:

(4) For all I know, *p* is true, but it’s not the case that the probability of *p* is greater than zero.<sup>xi</sup>

And (4) is clearly equivalent to

(5) For all I know, *p* is true, but the evidential probability of *p* is zero,

which in turn is equivalent to

(6) I don’t know that *p* is false, but the evidential probability of not-*p* is 1.

But if not-*p* is evidentially certain for me, then except in very strange cases I will know that *p*.

Thus on the semantics we propose (1)-(3) *are* in considerable tension; and in contexts where one believes according to the evidence<sup>xiii</sup> (and in unGettiered fashion) they’re likely contradictory, just as they would be on the “received view” of epistemic possibility. So we are at no disadvantage here. And our account of the relevant utterances’ infelicity is of a piece with our account of the typical oddity of CKAs: both boil down to speakers’ conveying conflicting information as to their degree of confidence in *p*, whether they regard not-*p* as a serious possibility or one that’s ruled out by their evidence,<sup>xiiii</sup> and so on (see I, above); that the usual conjuncts of CKAs don’t occur in (1)-(3) is beside the point.

### *III. On the sometime felicity of CKAs*

Dodd's third criticism is that by underwriting the felicity of CKAs<sup>xiv</sup> with a probabilistic (*evidentially* probabilistic) theory of epistemic possibility, we are not in fact establishing that they express truths. The pattern of the critique and our reply is the same here as it was with the first criticism: Dodd envisions us relying on a general principle which we neither assume nor need; he then illustrates the falsity of the principle with cases that leave untouched our treatment of CKAs, but which happen to fail on their own terms.

Dodd writes: "They seem to be assuming that the fact that a CKA can be uttered felicitously means it doesn't have a contradictory content" (p. 391). We assume no such thing.<sup>xv</sup> The dialectic is rather this: Stanley (2005) takes the infelicity of CKAs to be due to their having obviously contradictory contents (as they do, given the received view of epistemic possibility); we point out that *this* wrongly predicts the infelicity of certain CKAs the contents of which are (given the received view) equally semantically defective. Then, without attempting to derive the latter from the former, we give reason to believe that certain CKAs are both felicitous and non-contradictory – hence that they are not felicitous *merely* because they are non-literal, as felicitous contradictions tend to be.

Arguably, Dodd's own initial examples (*ibid.*) of felicitous contradictions are not particularly good, because it's not obvious that they're anything more than surface-contradictory. "I love you and I hate you," e.g., does not express two summary judgments; it expresses two separate judgments along two separate dimensions (e.g., *I love you in respect of your romantic nature* and *I hate you in respect of your philandering ways*). The other examples—"It was the best of times, it was the worst of times," "He's a genius but he's stupid," and "I'm happy and I'm sad"—are all of this nature, and so don't clearly illustrate that there could be felicitous contradictions.

Better are those examples of Dodd's which take our initial examples of (sometimes-)felicitous CKAs and substitute, in the first clause, claims about certainty for the claims about knowledge. There is, for example, "I'm sure they are going to lose, but I'm going to carry on watching just in case."<sup>xvi</sup> About such utterances, Dodd says: "To me all of these are just as felicitous as [Dougherty and Rysiew *et al.*'s examples]. However, I think they're clearly contradictory, or at least it's clear their speakers aren't committed to their literal truth" (p. 392). We're happy to allow the disjunction, but only because the second disjunct is very plausibly true. To illustrate, take Dodd's example: "I'm sure I have hands, even though it's possible I'm a handless brain in a vat." What we say about this depends on how we think about certainty (being "sure"). Here, we address the two possibilities that come naturally to mind:

A. Unger (1975), e.g., has famously argued that 'certain' is an *absolute* term, and that to be absolutely certain (hence, certain) is to be of the absolutely dogmatic attitude that no possible evidence could ever affect one's confidence as to *p*. But no one who was *literally certain* that they had hands could consistently express doubt about that proposition or take the possibility of their handlessness to have non-zero probability of their evidence. The fact is we use (absolute) certainty attributions non-literally all the time. One of the most common such uses is to avoid awkwardness in criticizing -- "I'm sure you've already thought of this, but what if the plane isn't on time?" In making such statements we're *never* sure in Unger's sense; in fact we're often *very unsure*, or indeed quite sure of the negation! This is a common form of non-literal speech and we all understand how it works (because we all use it!).

B. On the other hand, if "sure" just means 'really confident', or if "I'm sure" just means 'I have no real [serious] doubt', we *are* very often certain of things, including the falsity of various skeptical possibilities. This species of certainty is, by our fallibilistic lights, perfectly

compatible with some not-*p* possibility's having non-zero probability on one's evidence. In which case, the relevant utterance isn't a candidate for a felicitous contradiction, because it's not contradictory at all. What to say then, about any residual sense that *it is*? By now, the answer should be familiar: expressing the sort of confidence just described conflicts with the sort of *doubt* one would be naturally understood as conveying were one to go to the trouble of saying that some skeptical scenario is "possible" – for again (Section I), why say *that* unless one thinks there's some real chance (given one's evidence) that it obtains, and thus unless one isn't so sure that it doesn't?

In terms of these two options, we think that the situation with CKAs resembles *B*: the relevant utterances are non-contradictory and sometimes felicitous, with any infelicity surrounding them being merely pragmatic. Of course, an *A*-type treatment of CKAs is perfectly possible, with attributions of knowledge rather than certainty being construed non-literally. And when it comes to knowledge, there *are* such type-*A* philosophers – most notably, infallibilists. However, comparing the situation described in *A* with CKAs, epistemologists would be just about unanimous in agreeing that while we are hardly ever literally certain (in the relevant, i.e. option-*A*, sense) that something is the case, we often enough know things of which we're not sure (again, in the relevant sense). So there's nothing like the kind of pressure towards non-literality in the case of uncertain knowledge as there is in the case of certainty, *especially* when claims to the latter are accompanied by indicated doubt.<sup>xvii</sup> And such pressure as there *is* towards thinking that attributions of knowledge should be construed non-literally typically derives precisely from such observations as that CKAs, e.g., very often '*sound* contradictory' – in which case, it's thought, fallibilism can't be right (e.g., Lewis 1996); then attributions of knowledge, if felicitous, could *only* be non-literal (e.g., Unger 1975). But our account undercuts just such

arguments: the infelicity of CKAs can be explained, compatibly with both their literal truth and the truth of fallibilism.

#### IV. Conclusion

Dodd raises three objections to our account of CKAs. The first and third objections attribute to us assumptions we do not need, do not accept, and in one case openly reject. The second objection was that we couldn't satisfy Dodd's "second desideratum" of explaining the infelicity of a class of "for all I know" statements without resorting to the receive view of epistemic possibility and in an overall "unified" manner. We've met that challenge as well. Along the way, and more generally, we've clarified just what our treatment of CKAs and our arguments for it are, and what they aren't. We take ourselves to have shown that, as far Dodd's discussion goes, any confusion surrounding concessive knowledge attributions can't be laid at our doorstep.

#### References

- Carston, Robyn. 1988. Implicature, explicature, and truth-theoretic semantics. In *Mental Representations: The Interface Between Language and Reality*, R. M. Kempson ed. Cambridge: Cambridge University Press: 155-181.
- Dodd, D. 2010. Confusion about concessive knowledge attributions. *Synthese* 172.3: 381–396.
- Dougherty, T. Forthcoming. "Fallibilism." *Routledge Companion to Epistemology*.

Dougherty, T., & Rysiew, P. 2009. Fallibilism, epistemic possibility, and concessive knowledge attributions. *Philosophy and Phenomenological Research*, 78.1: 123-132.

Grice, H. P. 1961. The causal theory of perception. *The Aristotelian Society: Proceedings, Supplementary Volume 35*: 121-152.

\_\_\_ 1989. *Studies in the Way of Words*. Harvard University Press.

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

Jeffrey, R. 1992. *Probability and the Art of Judgment*. Cambridge 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.

\_\_\_ 2007. Speaking of knowing. *Noûs* 41.4: 627-662.

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

Unger, P. 1975. *Ignorance*. Oxford: Clarendon Press.

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

---

<sup>i</sup> Thanks to Patricia Blanchette and Jeffrey Speaks for helpful comments when these ideas were presented at Notre Dame. Thanks too to the faculty at Baylor University, and to two anonymous reviewers for *Synthese*.

<sup>ii</sup> “Confusion about concessive knowledge attributions,” this journal 172.3: 381–396. Below, parenthetical page numbers refer to Dodd’s paper.

<sup>iii</sup> Dodd (pp. 383; cf. 389 n. 12) states categorically that CKAs “are infelicitous in normal, non-philosophical contexts,” and less obviously so in philosophical ones only because there, our intuitions are “all over the place”. Using examples (Dougherty and Rysiew 2009, pp.130-131), we have defended them as perfectly upright citizens of conversation in certain everyday contexts -- not just (*pace* Dodd) when they concern “certain far out possibilities”.

<sup>iv</sup> We eschew “implicate” deliberately, avoiding “the compulsion to treat all pragmatically derived meaning as implicature” (Carston 1988, p. 176). In implicatures, properly so-called, one means what one says but something else besides; part of our treatment of CKAs involves raising the question of whether speakers do always mean what the relevant sentences (by our lights) semantically express, as opposed to what they thereby convey.

<sup>v</sup> Which of the latter is the best notion is a question that’s beyond the scope of this reply. We’ll use “non-zero” here for the sake of simplicity.

<sup>vi</sup> Using Jeffrey Conditioning, a proposition can have non-zero probability even if its negation is entailed by an item that is in one’s evidence. For example, suppose *The cab was green* gets in my evidence set in virtue of my having an experience as of a green cab. That proposition clearly entails the negation of *The cab was yellow* but since it might have been dark and hazy, my probability for the former proposition will be somewhat weak, say .85. Thus a yellow cab is still

---

an epistemic possibility for me. These sorts of considerations make Williamson's E = K thesis orthogonal to our discussion.

<sup>vii</sup> That is, this is what happens whenever there is cancelation. That is not to say that cancelation occurs whenever this pattern is followed. Nor should it be assumed that when a cancelation is signaled, it will always be recovered successfully by an interlocutor. In naturally-occurring CKAs (we describe some in our paper), the conversational setting (shared interests and information, etc.) makes it clear what *is* intended by the second conjunct of the CKA, where this is nothing that conflicts with the first conjunct: what is intended, rather, is the literal (and true) reading of that clause. However, where such stage-setting is absent -- as it is when philosophers narrow their focus to the bare, abstract, 'I know that *p*, but it's possible that not-*p*' -- it is entirely natural to 'hear' in the utterance of the second conjunct what's typically inferable therefrom -- that the speaker doesn't (take himself to) know that *p* (or that not-*p*), that the speaker isn't entirely confident that *p*, and so on. So, when it's mutually obvious that "*q*" usually conveys that not-*p*, the bare adding of "but *p*" (to "*q*") is as likely to occasion puzzlement as to why the speaker chose to express himself in *that* way as it is serve as an effective cancelation (compare Grice 1961, 137-138). Hence the fact, on our account, that CKAs can be odd-sounding even though they're very often literally true.

<sup>viii</sup> Depending on how one understands infelicity, we need to argue that a given CKA is either (a) true but infelicitous, the infelicity explaining the intuition of falsehood, or (b) that it is true and not even infelicitous (due to successful cancelation), with any persisting intuition of falsehood being explained by the fact that the one alleging falsehood is missing the fact of the cancelation (see next note).

---

<sup>ix</sup> Here is one place where it matters that we not treat all merely pragmatically conveyed information as an implicature (see n. 4): insofar as what's thus conveyed *is* what one means (versus, something one means in addition to what one says), its being conveyed, while cancelable, is liable to be less obviously comfortably canceled, especially given our tendency to read merely pragmatically-generated information onto the words uttered (see Rysiew 2001, 2007).

<sup>x</sup> Dodd makes ever-more-urgent reference to his footnote 11 (p. 388), which points out that Williamson's (2000) E=K thesis would collapse our evidence-based theory of epistemic possibility into the received knowledge-based view. We noted this ourselves (Dougherty and Rysiew 2009, p. 127, n. 5), though Dodd is wrong to suggest (p. 388) that it's in order to avoid this result that we reject E=K. Quite apart from any views on epistemic possibility and/or CKAs, we, like many others, take there to be abundant reason to reject Williamson's thesis. It's for independent reasons too that we now prefer the probabilistic version of our principle of epistemic possibility, in which case our two versions of our account might come apart (but it's not clear). The second version is rooted in the idea of epistemic probability. As noted above (n. 6), because of the generalized law of total probability --  $P(h) = \sum_i P(h/e_i)P(e_i)$  — a hypothesis can have less than unit probability *even if* E=K because each individual item of knowledge is itself less than certain. This assumes, with Jeffrey (1992), that observational data are less than certain, an idea of the essence of fallibilism (see Dougherty forthcoming).

<sup>xi</sup> “Definitely,” as in (1), is a modal auxiliary equivalent to “not possibly not.”

<sup>xii</sup> Recall that when one utters ‘*p*’ one represents oneself as believing that *p*.

<sup>xiii</sup> In his “Conclusion” (p. 394), Dodd suggests that maybe the conveying of *this* sort of information holds the key to a more promising fallibilism-friendly treatment of CKAs. As we've

---

seen, however, the conveying of this sort of information is already part of our own account of CKAs. (See too Rysiew 2007; that paper makes no explicit appeal to the present account of epistemic possibility, but it's perfectly consonant with both it and the current handling of CKAs.) Once again, Dodd seems to be working with an impoverished conception of just what's wrong with infelicitous CKAs on our view. (In the same way, Dodd says that Rysiew's 2001 treatment of CKAs has "the knowledge conjunct do all the pragmatic work" (*ibid.*). But that's not so either – see, e.g., Section 6 of that paper.)

<sup>xiv</sup> Dodd notes that we represent some cases from Hawthorne (2004, p. 24, n. 60; suggested by Weatherson and Tamar Gendler) as CKAs which are not technically so. While it's true that in those cases the one who says "I know that *p*" doesn't also assert "but possibly not-*p*", as Dodd himself allows they do act as if not-*p* is a genuine possibility and thus represent themselves as believing as much. (The idea of an act representing that one has certain mental states is a well-entrenched feature in the literature (e.g., Williamson 2000 and Unger 1975 before him).) But no matter, we like Dodd's own examples well enough, and as he points out there are others in the literature.

<sup>xv</sup> Nor, by the way, do we take CKAs to "provide evidence for fallibilism" (Dodd, p. 382): we think that CKAs express the fallibilistic idea, that CKAs are not semantically defective, that they're sometimes felicitously uttered, and so on. But it's nothing about CKAs that gets us thinking that fallibilism is correct.

<sup>xvi</sup> Arguably, the variant "There's no way they are going to lose, but I'm going to carry on watching just in case" is best read as "There is no *plausible* way...", in which case the problem doesn't arise here.

---

<sup>xvii</sup> For further discussion of problems facing such attempts to explain away the anti-sceptical appearances of our knowledge-attributing practices, even while much of the data one might glean from those practices probably shouldn't be taken at face value, see Rysiew 2007.