

COMMENTS ON W. HOLLIDAY'S  
"EPISTEMIC CLOSURE AND EPISTEMIC LOGIC I"

PAUL ÉGRÉ

**1. Overview.** In "Epistemic Closure and Epistemic Logic I", the first part of a tour-de-force inquiry into the semantics and logic of knowledge attributions, Holliday formalizes and compares several influential accounts of knowledge and its truth-conditions. Holliday's central result, the Closure Theorem, offers a rigorous and precise syntactic characterization of the demands made by different theories on epistemic closure. In so doing, Holliday achieves an unprecedented level of rigor and mathematical precision in spelling out the predictions of various theories relative to the problem of epistemic closure. Because the paper raises so many interesting issues, my comments will hardly do justice to the richness of its content and results. My goal here will be more limited, and will essentially concern two aspects of the problem of closure in relation to Holliday's account, namely: i) the question of the relative weights of subject factors and attributor factors in shaping the truth conditions of knowledge attributions, in particular for semantics based on relevant alternatives, and ii) the connection between the problem of epistemic closure and the problem of unattended possibilities in epistemic logic, in particular with regard to the trilemma that results from Holliday's Closure Theorem.

**2. Holliday's trilemma.** Before we go into these issues, a reminder of the philosophical underpinnings of the problem of epistemic closure will be useful. The problem of epistemic closure is fundamentally a version of the perennial problem of scepticism. I know that I have hands. My having hands implies that I am not a brain in a vat, artfully deceived by an evil scientist or demon into thinking that I have hands. But I do not know, the sceptic points out, that I am not such a brain in a vat. If knowledge is closed under implication, we have a problem. Relying on the force of his observation, the sceptic concludes, by modus tollens, that I do not know that I have hands after all. Holding on to the strength of the main premise, the dogmatist counters, by modus ponens, that I do know that I am not a brain in a vat after all. For Dretske or Nozick, a third way is to maintain with the dogmatist that I know I have hands, and to deny with the sceptic that I know I am not a brain in a vat, but to resist that knowledge is closed under implication.

The rejection of epistemic closure is more costly than we might have thought at first, and the value of Holliday's contribution is to give an exact assessment of this cost. One way in which Holliday's result can be interpreted is as posing a dilemma, or rather, a trilemma, seeing that the second horn of the former subdivides: either we deny epistemic closure, like Dretske or Nozick, but then we end up denying further closure properties of knowledge that seem intuitively desirable (such as the principle  $K(\phi \wedge \psi) \rightarrow K\phi$ ); or we maintain epistemic closure, but then we must confront a subordinate dilemma: either we

preserve closure by sticking to invariantism about knowledge (Holliday’s C-semantics in the paper), thereby incurring the risk of scepticism (or dually, of dogmatism); or we adopt a contextualist semantics, of the kind suggested by Lewis (Holliday’s L-semantics), but then we run the risk of making knowledge too easily attainable (what Holliday calls the problem of vacuous knowledge, namely knowledge acquired cheaply by suitably restricting the space of relevant alternatives).

Faced with the trilemma, Holliday’s inclination appears to be to favor contextualism about knowledge. I agree that this is a reasonable choice, but before we can assess whether it is the best choice, I think we need to know how compelling the trilemma is, and how much the trilemma depends on the particular framework chosen by Holliday to represent the various theories under discussion.

**3. Subject vs. attributor.** The first aspect I would like to question in Holliday’s approach concerns the relative weight of subject factors and attributor factors in the truth conditions of knowledge ascriptions. A very important observation about knowledge attributions is that they involve both a subject about whom we make ascriptions, and an attributor who ascribes knowledge (see Holliday’s emphasis on the distinction pg. 9). Knowledge is a function of at least those two perspectives. Prima facie, both perspectives appear to be reflected in Holliday’s relevant alternative models, with the distinction between the accessibility relation (whether two worlds are equivalent or not) and the preorder on worlds (whether one world is more relevant than another or not). However, Holliday warns us that: “These models represent the epistemic state of an agent from a third-person perspective. We should not assume that anything in the model is something that the agent has in mind”. This is fine, but right away he adds that “Just as the model is not something that the agent has in mind, it is not something that particular speakers attributing knowledge to the agent have in mind either.”.

I agree with Holliday that relevance in particular can be both a subject factor and an attributor factor, but I end up being confused on which feature of the model represents which perspective. As I see it, relevance is better handled *primarily* as an attributor’s factor, simply to reflect the fact that the attributor (who may be an ideal referee, or just some situated agent in a particular community) is the ultimate judge in fixing what counts as pragmatically relevant in ascribing knowledge. An abstract way of articulating a relevant alternative semantics would thus be by distinguishing between the subject’s discriminative capacities relative to the whole set of worlds, and the attributor’s assessment of which alternatives are to be considered. In principle, we should probably duplicate perspectives completely, and have, for each subject and attributor, both a respective discrimination set and a respective relevancy set. But one way to keep matters more simple on a first approximation is to put relevance wholly on the side of the attributor, and discrimination on the side of the subject. In particular, the most basic relevant alternative semantics for knowledge ascriptions can be stated simply by distinguishing, relative to a world  $w$  in a set  $W$ , the set  $R_a(w)$  of the possibilities that the attributor  $a$  thinks the subject  $s$  should entertain, and the set  $D_s(w)$  of possibilities that the subject  $s$  cannot discriminate from  $w$ , irrespective of whether  $s$  actually entertains those or not:

- (1)  $w \models K_s \phi$  iff for every  $w' \in R_a(w)$ : if  $w D_s w'$ , then  $w' \models \phi$ .

This semantics preserves the principle of epistemic closure, just like the Lewisian L-semantics in Holliday's account, and it can account for the same contrast cases, like the medical diagnosis example discussed by Holliday.<sup>1</sup> Of course, the L-semantics can be seen as one way of specifying this abstract semantics, letting  $R_a(w)$  correspond to  $Min_{\leq w}(W)$ , and  $D_s$  be the arrow relation in Holliday's model. There are two differences however between those ways of articulating a Lewisian semantics: one is the fact that we have attached relevance to the attributor's perspective, and discrimination to the subject's perspective; the second is that Holliday's treatment of relevance gives the notion more structure from the start. In Holliday's framework, alternatives can be *more or less* relevant relative to a world, whereas in the simple-minded semantics presented above, an alternative is simply relevant or not.

**4. Attention.** This brings me to the second aspect I would like to question. The sceptical argument, remember, is that: from  $K_s \phi$  (I know I have hands), and  $K_s(\phi \rightarrow \neg\psi)$  (I know that if I have hands, then I am not a brain in a vat), we must infer  $K_s \neg\psi$  (I know I am not a brain in a vat), even for a  $\psi$  that is a far-fetched possibility in comparison to  $\phi$ . As talk of far-fetched possibilities suggests, it is quite natural to introduce an ordering on possibilities to capture the difference between  $\phi$  and  $\psi$ , the way Holliday does. If we use the semantics laid out above, we have no way of capturing this difference. Nevertheless, I think we can do something similar by distinguishing a second subject factor beside discrimination, namely *attention* to possibilities. The idea would be that, when I claim to know that I have hands, I am failing to attend to the possibility that I am a brain in a vat. Because of that, I may also fail to attend to the implication that, if I have hand, then I am not a brain in a vat. In other words, instead of rejecting epistemic closure, the proposed solution to the sceptical paradox would consist in accepting the validity of closure, but in denying systematic knowledge of the relevant conditional, for cases in which the conditional mentions a possibility that is unattended and properly ignored (see Lewis 1996).

In Aloni, Egré and de Jager (2009), truth conditions for knowledge ascriptions are proposed along those lines, building on de Jager's work on unawareness (see de Jager 2009). We define an attention and relevance model (for a single subject  $s$  and attributor  $a$ ) as a structure  $(W, D_s, R_a, E_s, A_s, V)$ , where  $W$  is the set of possibilities,  $D_s$  is the discrimination relation over  $W$ ,  $R_a$  the relevance relation,  $E_s$  a function that to each world associates the possibilities the agents entertains, and  $A_s$  a function that to each world associates the sentences that the agent is aware of.  $E_s$  and  $A_s$  together serve to define the notion of attention to a possibility. Knowledge ascriptions are evaluated as follows, letting  $\mathcal{L}(X)$  be the fragment of the whole language consisting of only atoms in  $X$ :

- (2)  $w \models K_s \phi$  iff  $\phi \in \mathcal{L}(A_s(w))$  and  $\forall w' \in E_s(w) \cup R_a(w)$ : if  $w' \in D_s(w)$ , then  $w' \models \phi$

---

<sup>1</sup>Let  $R_a(w)$ , the alternatives the professor considers relevant: in the first context, only include the alternative  $w_c$ : both  $A$  and  $B$  would know  $c$ ; in a second context, let it include also  $w_{c'}$ , then only  $A$  knows  $c$ ; in a third, if  $R_a$  also includes alternative  $w_x$ ,  $A$  does not know  $c$  any more.

This says that  $\phi$  is known provided the agent is aware of  $\phi$  and  $\phi$  holds in every alternative that is entertained and in every alternative that is considered relevant by the attributor. Now, the important point is that  $K_s\phi$  and  $K_s(\phi \rightarrow \psi)$  will entail  $K_s\phi$  under suitable provisos on  $\mathcal{L}(A_s)$ . This is plausible enough, if we consider that, once I am attending to the proposition that, if I have hands, then I am not a brain in a vat, then I am also attending to the proposition that I am not a brain in a vat. However, it will be possible to have a model in which:  $K_s\phi$  is true, but in which both  $K_s(\phi \rightarrow \neg\psi)$  and  $K_s(\neg\psi)$  are false, if indeed  $\psi$  is a possibility that is “properly ignored” by the agent, namely both irrelevant and not attended to. On that view of the sceptical scenario, epistemic closure need not be the culprit (the principle can be valid, but need not be sound). Importantly, this solution hinges on the fact that an implication  $\phi \rightarrow \psi$  can be universally true over  $W$  without necessarily being attended to, and so, without being automatically known (in contrast to what a plain L-semantics or a basic semantics like the one in (1) would predict). This means that I can fail to know that I am not a brain in a vat, not necessarily for lack of evidence, but for failing to attend to the possibility.

To conclude on this, I wonder if the notion of unattended possibility, in Holliday’s relevant alternative models, could be captured in terms of a possibility that is so remote in the ordering that it is tantamount to being unattended. Or whether inattentiveness would have to be captured along a distinct dimension.

#### REFERENCES

- [1] Aloni M., Egré P. and de Jager, T. (2009). Knowing whether A or B. *Synthese*. DOI 10.1007/s11229-009-9646-1.
- [2] De Jager T. (2009). ‘Now that you mention it, I wonder...’: Awareness, Attention, Assumption. PhD dissertation, ILLC, Amsterdam.
- [3] Holliday W. (2013). Epistemic Closure and Epistemic Logic I: Relevant Alternatives and Subjunctivism. *Journal of Philosophical Logic*.
- [4] Lewis D. (1996). Elusive Knowledge. *Australasian Journal of Philosophy* 74(4):549-567