Response to Egré and Xu[*]

Wesley H. Holliday

University of California, Berkeley
wesholliday@berkeley.edu

Abstract. In this note, I respond to comments by Paul Egré and Xu Zhaoqing on my "Epistemic Closure and Epistemic Logic I: Relevant Alternatives and Subjunctivism" (*Journal of Philosophical Logic*).

I want to begin by thanking Paul Egré and Xu Zhaoqing for their thoughtful and thought-provoking remarks on my paper, "Epistemic Closure and Epistemic Logic I: Relevant Alternatives and Subjunctivism" (Holliday 2013a, hereafter 'EC&ELI'). It is a privilege to have one's work receive close readings and constructive comments from international colleagues. In the space available here, I will try to address some of the main points raised by Egré and Xu. Throughout I will presuppose familiarity with the summary of EC&ELI in this volume.

Response to Paul Egré

As Egré observes, EC&ELI sets up a trilemma. Each of the theories of knowledge considered in the paper encounters one of the three horns: the Problem of Skepticism (C-semantics), the Problem of Vacuous Knowledge (L/S-semantics), and the Problem of Containment (D/H/N/S-semantics). In EC&ELI, I do not offer a solution to this trilemma, but rather try to systematically investigate the third horn, the Problem of Containment. Given how little I let on in EC&ELI about my own response to the trilemma, it is not unreasonable that Egré infers that "Faced with the trilemma, Holliday's inclination appears to be to favor contextualism" about knowledge attributions (see Holliday 2012a for more on contextualism). However, for the record I should say that I do not think that contextualism about knowledge attributions is the key to solving the trilemma. Instead, I think the key to solving the trilemma is to replace what I call the Standard Alternatives Picture, in which all of the theories of knowledge in EC&ELI fit as special cases, with a new and improved Multipath Picture of Knowledge [Holliday, 2013b,c, 2012b]. The Multipath Picture of Knowledge is compatible with contextualism; but it is this new picture of *knowledge*, not contextualism about *knowledge attributions*, that in my view solves the trilemma. All of this is discussed in Holliday 2013b,c, which developed out of Holliday 2012b.

With that clarification made, let me turn to Egré's two main points:

---

[*] Published in *Logic Across the University: Foundations and Applications* (College Publications), eds. Johan van Benthem and Fenrong Liu, 39-46.

Subject vs. Attributor. Egré helpfully emphasizes the need to be clear about which aspects of the relevant alternatives (RA) models in EC&ELI depend on the *knowing subject* and which depend on the *person attributing knowledge* to the subject. Egré sketches what he considers a more basic and abstract relevant alternative semantics, which makes explicit a contribution of the subject and a contribution of the attributor. Formally, what Egré labels as (1) is equivalent to the L-semantics of EC&ELI.[1] However, there is a conceptual difference.

In RA models $\mathcal{M} = \langle W, \rightarrow, \preceq, V \rangle$ for L-semantics, the $\rightarrow$ relation represents what possibilities the subject (not the attributor) can discriminate between, so there is no conceptual difference here between L-semantics and the semantics Egré sketches. However, there is a conceptual difference elsewhere: when I presented RA models for L-semantics in EC&ELI, I said that the set $\text{Min}_{\preceq_w}(W)$ of *relevant* worlds at w can depend not only on "attributor factors" but also on "subject factors" (see DeRose 2009, 30f), whereas Egré says that the set of relevant worlds at w, written as $R(w)$ instead of $\text{Min}_{\preceq_w}(W)$,[2] can only depend on *attributor* factors. (In this sense, Egré's interpretation is less general.)

According to Egré, "relevance is better handled *primarily* as an attributor's factor, simply to reflect the fact that the attributor …is the ultimate judge in fixing what counts as pragmatically relevant in ascribing knowledge," so "a first approximation is to put relevance wholly on the side of the attributor, and discrimination on the side of the subject." One can certainly interpret the relevance relations in RA models as depending only on the attributor, but doing so would be highly controversial. According to Dretske [1981, 2004], for example, what possibilities are relevant does not depend at all on the conversational context of the attributor, but rather on objective features of the subject's environment, of which both the subject and attributor might be unaware. In one of his examples, Dretske [1981] considers whether a birdwatcher in Wisconsin must eliminate the possibility that the bird he sees on a lake is a Siberian Grebe in order to know that it is a Gadwall. (We are to suppose that a Gadwall cannot be distinguished from a Siberian Grebe unless one sees the underbelly of the bird in flight.) Roughly, Dretske suggests that if, as a matter of fact, there are many Siberian Grebes in the birdwatcher's area, then the birdwatcher needs to eliminate the possibility that what he sees is a Siberian Grebe in order for an attributor to truly say of him that he knows the bird is a Gadwall—in that sense, it is a relevant possibility—*even if the attributor is unaware of Siberian*

---

[1] This assumes that Egré requires $w \in R_a(w)$, as required for $K\varphi \rightarrow \varphi$ to be valid. However, since Egré takes $R_a(w)$ to be the set of "possibilities that the attributor a thinks the subject s should entertain," it is unclear whether this understanding guarantees that $w \in R_a(w)$ always holds. After all, an attributor in w might not know that w is the actual world and therefore might not think that w is one of the possibilities that s should entertain. Perhaps Egré's idea is that every attributor thinks that the actual world (*de dicto*), whatever it is, should be entertained?

[2] Note that for L-semantics, the rest of the relevance relation $\preceq_w$ beyond $\text{Min}_{\preceq_w}(W)$ does not matter in a fixed context, so there is no substantive difference between $R(w)$ and $\text{Min}_{\preceq_w}(W)$. Yet when it comes to the dynamics of *context change*, the rest of the relevance relation $\preceq_w$ may indeed matter, as discussed in Holliday 2012a.

Grebes. Most contextualists also grant that possibilities may be relevant in virtue of these kinds of "subject factors" of which attributors may be unaware. As I say in EC&ELI, "possibilities may be relevant and hence should be included in our model, even if the attributors are not considering them (see DeRose 2009, 33)."

In any case, I think Egré makes an excellent point about keeping explicit track of the subjects and attributors, who remained nameless in EC&ELI. Since multiple attributors may share the same context, let us keep track of subjects (agents) and contexts (possibly associated with groups of agents). Given a set $\mathcal{S}$ of agent symbols and a set $\mathfrak{C}$ of context symbols, define an extended language:

$$\varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid K_s^{\mathfrak{c}}\varphi$$

with $s \in \mathcal{S}$ and $\mathfrak{c} \in \mathfrak{C}$, reading '$K_s^{\mathfrak{c}}\varphi$' as "agent s counts as knowing φ relative to context $\mathfrak{c}$." The generalization of RA models for the extended language is just what one would expect: a *multi-agent and multi-context RA model* is a tuple

$$\mathcal{M} = \langle W, \{\twoheadrightarrow_s\}_{s\in\mathcal{S}}, \{\preceq^{s,\mathfrak{c}}\}_{s\in\mathcal{S},\, \mathfrak{c}\in\mathfrak{C}}, V \rangle,$$

where for each $s \in \mathcal{S}$ and $\mathfrak{c} \in \mathfrak{C}$, $\langle W, \twoheadrightarrow_s, \preceq^{s,\mathfrak{c}}, V \rangle$ is an RA model as in EC&ELI. The generalization of C/D/L-semantics from EC&ELI to multi-RA models is also straightforward. For a well-founded[3] multi-RA model $\mathcal{M}$, define $\mathcal{M}, w \vDash_x \varphi$ and $[\![\varphi]\!]_x^{\mathcal{M}} = \{v \in W \mid \mathcal{M}, v \vDash_x \varphi\}$ as follows (for notation, see EC&ELI):

$$\mathcal{M}, w \vDash_c K_s^{\mathfrak{c}}\varphi \quad \text{iff} \quad \forall v \in \overline{[\![\varphi]\!]}_c^{\mathcal{M}} : w \not\twoheadrightarrow_s v;$$
$$\mathcal{M}, w \vDash_d K_s^{\mathfrak{c}}\varphi \quad \text{iff} \quad \forall v \in \text{Min}_{\preceq_w^{s,\mathfrak{c}}}\big(\overline{[\![\varphi]\!]}_d^{\mathcal{M}}\big) : w \not\twoheadrightarrow_s v;$$
$$\mathcal{M}, w \vDash_l K_s^{\mathfrak{c}}\varphi \quad \text{iff} \quad \forall v \in \text{Min}_{\preceq_w^{s,\mathfrak{c}}}(W) \cap \overline{[\![\varphi]\!]}_l^{\mathcal{M}} : w \not\twoheadrightarrow_s v.$$

One can generalize CB models and H/N/S-semantics in the analogous way.

Generalizing to multiple agents and multiple contexts raises interesting new issues, supporting Egré's call to make these aspects of the model explicit:

With multiple agents, we can consider multi-agent closure principles. For example, one can check that the principle $K_s K_u \varphi \to K_s \varphi$ is not valid according to D/H/N/S-semantics; thus, e.g., according to Nozick's tracking theory, an agent s can know that an agent u knows φ without s knowing φ herself.

With multiple contexts, we can consider cross-context knowledge, as in the so-called Problem of Factivity for contextualism [Williamson, 2001, Wright, 2005, Brendel, 2005, Baumann, 2008], which I cannot resist repeating here:

According to standard contextualism, as a result of a skeptic's argumentation a contextualist u might find herself in a context $\mathfrak{s}$ relative to which she does not count as knowing ordinary propositions. In such a case, the contextualist might like to respond to the skeptic by claiming that she still counts as knowing those ordinary propositions relative to an ordinary context $\mathfrak{o}$. However, by making such a claim in context $\mathfrak{s}$, she would be claiming something that is *impossible for her to know* in $\mathfrak{s}$, as a simple argument shows. Our initial assumption was $\neg K_u^{\mathfrak{s}} p$. Suppose for reductio that $K_u^{\mathfrak{s}} K_u^{\mathfrak{o}} p$. Relative to any context, knowledge is

---

[3] I.e., in which each preorder $\preceq_w^{s,\mathfrak{c}}$ is well-founded as explained in EC&ELI.

factive, reflected by the validity of $K_u^{\mathfrak{o}}p \to p$ according to all of the semantics; and we can assume that the contextualist knows the factivity principle relative to any context, reflected by the validity of $K_u^{\mathfrak{s}}(K_u^{\mathfrak{o}}p \to p)$ according to all of the semantics. Finally, many contextualists claim that knowledge is closed under known implication relative to any fixed context, reflected by the validity in L-*semantics* of $(K_u^{\mathfrak{c}}\varphi \wedge K_u^{\mathfrak{c}}(\varphi \to \psi)) \to K_u^{\mathfrak{c}}\varphi$, an instance of which is

$$(K_u^{\mathfrak{s}}K_u^{\mathfrak{o}}p \wedge K_u^{\mathfrak{s}}(K_u^{\mathfrak{o}}p \to p)) \to K_u^{\mathfrak{s}}p.$$

Putting it all together, we derive $K_u^{\mathfrak{s}}p$, reflected by the validity of $K_u^{\mathfrak{s}}K_u^{\mathfrak{o}}\varphi \to K_u^{\mathfrak{s}}\varphi$ according to L-semantics. But $K_u^{\mathfrak{s}}p$ contradicts our initial assumption of $\neg K_u^{\mathfrak{s}}p$; hence $K_u^{\mathfrak{s}}K_u^{\mathfrak{o}}p$ is impossible. Moreover, it is plausible that the contextualist can follow this derivation and come to know that $K_u^{\mathfrak{s}}K_u^{\mathfrak{o}}p$ is impossible. Hence a weak *norm of assertion*, namely that you should not assert something (in a context) that you know is impossible for you to know (in that context), prohibits the contextualist from replying to the skeptic as she might have liked to in $\mathfrak{s}$. Without going further into this Problem of Factivity and related issues, the argument above suffices to show the interest of the multi-context generalization.

Attention. Egré's second main point concerns the notion of attention to possibilities (and awareness of sentences). Rather than trying to summarizing Egré's nuanced discussion of attention, I will go directly to my three responses.

First, insofar as an attributor or subject paying attention to a possibility would affect the relevance of that possibility, it would affect the relevance relation $\preceq_w$ in RA models. For example, according to Lewis's [1996] *Rule of Attention*, any possibility that the attributors are attending to in a context is relevant in that context, i.e., in $\mathrm{Min}_{\preceq_w}(W)$.[4] On this view, if the attributors attend to additional possibilities, then $\mathrm{Min}_{\preceq_w}(W)$ will enlarge. (For formal modeling of the dynamics of context change, see Holliday 2012a.) One might also hold that a subject s attending to a possibility tends to make that possibility relevant according to $\preceq_w^s$ (whether or not the attributors are attending to it). Indeed, a whole range of views about the relation between attention and relevance are possible and compatible with modeling relevance as in RA models.[5]

Second, on Egré's proposed solution to the skeptical problem, I do not think it is sufficient to solve the problem to say that agents are not paying attention to—and thus don't know—conditionals of the form $p \to \neg SK$ where p is an ordinary proposition and SK a skeptical hypothesis. It is true that if an agent does not know $p \to \neg SK$, this blocks *one* skeptical argument for $\neg Kp$, namely the one using $\neg K\neg SK$, $(Kp \wedge K(p \to \neg SK)) \to K\neg SK$, and modus tollens. But the

---

[4] Although for L-semantics all that matters is whether a possibility is in or out of $\mathrm{Min}_{\preceq_w}(W)$, in D-semantics the rest of the relation $\preceq_w$ matters, and a D-semantic contextualist could take the degree of attention that attributors are paying to a possibility to affect its ranking in the relevance ordering $\preceq_w$.

[5] I should remind the reader that my own view of knowledge rejects the simple world-ordering picture used in RA models for D- and L-semantics (see Holliday 2013b,c, 2012b), but I am trying to motivate that picture here on its own terms.

essence of the skeptical challenge is this: if an agent cannot eliminate the kind of possibilities described by skeptical hypotheses, how can she know ordinary propositions about the world, given that the skeptical hypotheses are hypotheses according to which she is *radically deceived* about the world? It does not seem to help here to plead ignorance, on the agent's behalf, of some conditionals. If anything, that makes the agent's epistemic situation look even *worse* according to the skeptic, who would otherwise be willing to grant that although the agent doesn't know much at all, at least she knows some obvious conditionals. Of course, there is much more to be said about skepticism (see, e.g., Stroud 1984, DeRose 1995), but there is no room to do so here (see Holliday 2013b).

Finally, Egré mentions a fascinating paper, Aloni et al. 2009, that I also cannot discuss here. But I will make one point about the truth clause for knowledge attributions, inspired by Aloni et al. 2009, that Egré labels as (2) in his comments. According to (2), for an attributor a to truly attribute knowledge of φ to a subject s, it is only required that s has eliminated the $\neg\varphi$-possibilities in the set $E_s(w) \cup R_a(w)$ of possibilities "entertained" by s or "considered relevant" by a. I agree with Dretske, however, that just because a subject isn't entertaining a $\neg\varphi$-possibility and an attributor doesn't consider it relevant, it doesn't follow that the subject need not eliminate that possibility for the attributor to truly attribute knowledge of φ to the subject; consider, e.g., a case in which a patient presents symptoms compatible with several common diseases, but her incompetent medical intern (the subject) and incompetent doctor (attributing knowledge to the intern) only think of one of them and fail to consider the others.

According to Egré's (2) and (1), the fewer possibilities the subject entertains and the attributor considers, the more the attributor can truly say that the subject knows. But just because a close-minded subject and carefree attributor do not entertain or consider possibilities that they *should* does not mean that the attributor can truly attribute knowledge to the subject thanks to their carefreeness and close-mindedness.[6] Knowledge, unlike mere opinion or belief, is an achievement that cannot be acquired so cheaply. A subject's entertainings and an attributor's considerings do not completely determine what counts as relevant; objective aspects of the world, such as the real frequencies of diseases, also affect what possibilities are relevant and must be eliminated for knowledge.

Response to Xu Zhaoqing

Xu helpfully offers three main comments about EC&ELI, as well as a number of interesting side notes. For reasons of space, I focus on the big-picture points.

---

[6] The formulation in Aloni et al. 2009 replaces $E_s(w) \cup R_a(w)$ in (2) with $E(w) \cup S(w)$, where $E(w)$ is the set of possibilities that *are* entertained by the subject and $S(w)$ is the set of possibilities that *should* be entertained by the subject. But it seems to me that there can be possibilities that need to be eliminated for knowledge but that neither are entertained nor *should* be *entertained* by the subject—since "entertaining" possibilities sounds like an intellectual act that takes time, while the agent's perception may quickly eliminate possibilities without so much thought.

Knowledge and Belief. Xu's first point concerns adding *belief* to RA models. In EC&ELI, for simplicity I did not represent an agent's belief in RA models separately from her knowledge (though I did in CB models). To justify this choice, I made two points about belief in relation to RA models. To set them up, suppose that we define *relevant alternatives and belief* (RAB) *models* as tuples $\mathcal{M} = \langle W, \rightarrowtail, D, \preceq, V \rangle$, where $\langle W, \rightarrowtail, \preceq, V \rangle$ is an RA model and D is a doxastic accessibility relation as in CB models, and redefine the truth clauses as follows:

$$\mathcal{M}, w \vDash_x B\varphi \quad \text{iff} \quad \forall v \in W \colon \text{if } wDv \text{ then } \mathcal{M}, v \vDash_x \varphi;$$
$$\mathcal{M}, w \vDash_c K\varphi \quad \text{iff} \quad \mathcal{M}, w \vDash_c B\varphi \text{ and } \forall v \in \overline{\llbracket \varphi \rrbracket}_c^{\mathcal{M}} \colon w \not\rightarrowtail v;$$
$$\mathcal{M}, w \vDash_d K\varphi \quad \text{iff} \quad \mathcal{M}, w \vDash_d B\varphi \text{ and } \forall v \in \mathrm{Min}_{\preceq_w}\big(\overline{\llbracket \varphi \rrbracket}_d^{\mathcal{M}}\big) \colon w \not\rightarrowtail v;$$
$$\mathcal{M}, w \vDash_l K\varphi \quad \text{iff} \quad \mathcal{M}, w \vDash_l B\varphi \text{ and } \forall v \in \mathrm{Min}_{\preceq_w}(W) \cap \overline{\llbracket \varphi \rrbracket}_l^{\mathcal{M}} \colon w \not\rightarrowtail v.$$

My first point about this in EC&ELI was that the switch from RA to RAB models would not change any of the main results of the paper; the Closure Theorem and the completeness theorem would easily extend to RAB models. My second point was that if we imposed the following constraint on RAB models, then the extra $B\varphi$ requirement for $K\varphi$ would become redundant anyway:

$$\forall w, v \in W : wDv \Rightarrow [v \in \mathrm{Min}_{\preceq_w}(W) \text{ and } w \rightarrowtail v],$$

which says that if v is compatible with what the agent believes in w, then v is relevant and uneliminated for the agent in w, ensuring the validity of $K\varphi \rightarrow B\varphi$. I will not discuss the plausibility of this constraint, which is a variant of Lewis's [1996] *Rule of Belief.* Instead, I will simply say that like Xu, I am curious about the question he raises: are there plausible constraints relating D, $\preceq$, and $\rightarrowtail$?

Existential Quantification and Neighborhood Models. Noting that truth clauses for $K\varphi$ usually use universal quantification, Xu points out that in connection with closure it would be interesting to consider truth clauses for $K\varphi$ that involve existential quantification, as in the case of neighborhood semantics, where $K\varphi$ is true at w iff there is a proposition in $N(w)$ that is (a subset of) the set of worlds where φ is true. My short response to Xu's point here is that according to my own positive view of knowledge, the Multipath Picture of Knowledge [Holliday, 2013b,c, 2012b] noted above, the truth clause for $K\varphi$ does have an existential character and the models for knowledge are indeed related to neighborhood models. So my response to Xu's point about existentials and neighborhoods is 'yes' and 'yes'! For the details, see Holliday 2013b,c, 2012b.

Stronger Implications. Finally, Xu raises the issue of stating closure principles with stronger kinds of implication than material implication. In EC&ELI, I noted that not only closure under known implication but also closure under known strict implication, $(K\varphi \wedge K\Box(\varphi \rightarrow \psi)) \rightarrow K\psi$, is invalid according to D/H/N/S-semantics. (Note, by contrast, that closure under strict equivalence, $(K\varphi \wedge K\Box(\varphi \leftrightarrow \psi)) \rightarrow K\psi$ is D/H/N/S-valid.) Xu asks whether if we interpret

a *counterfactual* $\Box\!\!\rightarrow$ using the same preorder used for K, the closure principle $(K\varphi \wedge K(\varphi \;\Box\!\!\rightarrow\; \psi)) \to K\psi$ is valid.[7] The answer is 'yes' for C/L-semantics and 'no' for D/H/N/S-semantics; the 'no' answer follows from the invalidity of $(K\varphi \wedge K\Box(\varphi \to \psi)) \to K\psi$, since the strict implication is stronger than the counterfactual, but to see the reason for D/H/N-semantics, observe that from the facts that the agent has eliminated the minimal $\neg\varphi$-worlds ($K\varphi$) and that the minimal φ-worlds are ψ-worlds ($\varphi \;\Box\!\!\rightarrow\; \psi$), it does not follow that the agent has eliminated the minimal $\neg\psi$-worlds ($K\psi$). Xu's question underscores an interesting open problem: axiomatize D/H/N/S-semantics for extended languages.

Conclusion

I am grateful to Egré and Xu for raising so many stimulating points, more than I could cover here. I hope, however, that some of the answers—or at least more good questions—arise out of the work that follows EC&ELI [Holliday, 2013b,c].

---

[7] I do not think a relation $\preceq_w$ for RA theories should be identified with an ordering for counterfactuals, so I take Xu's question to be mainly about subjunctivist theories.

Bibliography

M. Aloni, P. Egré, and T. de Jager. Knowing whether A or B. *Synthese*, 2009. DOI 10.1007/s11229-009-9646-1.

P. Baumann. Contextualism and the Factivity Problem. *Philosophy and Phenomenological Research*, 76:580–602, 2008.

E. Brendel. Why contextualists cannot know they are right: self-refuting implications of contextualism. *Acta Analytica*, 20(2):38–55, 2005.

K. DeRose. Solving the Skeptical Problem. *The Philosophical Review*, 104(1): 1–52, 1995.

K. DeRose. *The Case for Contextualism: Knowledge, Skepticism, and Context.* Oxford University Press, New York, 2009.

F. Dretske. The Pragmatic Dimension of Knowledge. *Philosophical Studies*, 40 (3):363–378, 1981.

F. Dretske. Externalism and Modest Contextualism. *Erkenntnis*, 61:173–186, 2004.

W. H. Holliday. Epistemic Logic, Relevant Alternatives, and the Dynamics of Context. In D. Lassiter and M. Slavkovik, editors, *New Directions in Logic, Language and Computation*, volume 7415 of *Lecture Notes in Computer Science*, pages 109–129. 2012a.

W. H. Holliday. *Knowing What Follows: Epistemic Closure and Epistemic Logic.* PhD thesis, Stanford University, 2012b. Revised Version, ILLC Dissertation Series.

W. H. Holliday. Epistemic Closure and Epistemic Logic I: Relevant Alternatives and Subjunctivism. *Journal of Philosophical Logic*, 2013a. Forthcoming.

W. H. Holliday. Fallibilism and Multiple Paths to Knowledge. *Oxford Studies in Epistemology*, 5, 2013b. Forthcoming.

W. H. Holliday. Epistemic Closure and Epistemic Logic II: A New Framework for Fallibilism. 2013c. Manuscript.

D. Lewis. Elusive Knowledge. *Australasian Journal of Philosophy*, 74(4):549–567, 1996.

B. Stroud. *The Significance of Philosophical Scepticism.* Oxford University Press, New York, 1984.

T. Williamson. Comments on Michael Williams' "Contextualism, Externalism, and Epistemic Standards". *Philosophical Studies*, 103(1):25–33, 2001.

C. Wright. Contextualism and Scepticism: Even-Handedness, Factivity, and Surreptitiously Raising Standards. *The Philosophical Quarterly*, 55(219):236–262, 2005.