Epistemic Closure and Epistemic Logic I: Relevant Alternatives and Subjunctivism. A Summary[⋆]

Wesley H. Holliday

University of California, Berkeley
wesholliday@berkeley.edu

Abstract. I summarize the formal framework and some of the main formal results in my "Epistemic Closure and Epistemic Logic I: Relevant Alternatives and Subjunctivism" (*Journal of Philosophical Logic*), adding a few methodological comments on formalization in philosophy.

Some topics in philosophy seem to cry out for formal analysis. The topic of *epistemic closure* in epistemology strikes me as a prime example. Its very name invokes the mathematical idea of a set being closed under a relation. In one of its simplest forms, the question is whether the set of *propositions known by an ideal logician* L must be closed under the relation of *logical consequence* from multiple premises: if $P_1, \dots, P_n$ are in the set of propositions known by L, and P_{n+1} is a logical consequence of $\{P_1, \dots, P_n\}$ (in which case L has at least deduced and come to *believe* P_{n+1}), must P_{n+1} be in the set of propositions known by L? Epistemic closure questions like this have lead to "one of the most significant disputes in epistemology over the last forty years" [Kvanvig, 2006, 256], as different philosophical theories of knowledge give different answers. Some theories support full closure, while others only support weaker closure principles.

One of the goals of my "Epistemic Closure and Epistemic Logic I: Relevant Alternatives and Subjunctivism" (Holliday 2013b, hereafter 'EC&ELI') was to determine exactly the extent of epistemic closure according to a family of standard theories of knowledge. The theories at the center of the closure debate, versions of the *relevant alternatives* (RA) and *subjunctivist* theories of knowledge, are often presented in a kind of modal picture, distinguishing relevant epistemic alternatives or close counterfactual possibilities from the rest of logical space. Thus, it seemed that a natural way to study the properties of these theories was to formalize them with some kind of models for epistemic modal logic. This was the path followed in EC&ELI, which I will sketch in this summary.

In my view, an application of formal methods in philosophy gains value if:

(1) *it is faithful to the philosophical views being formalized*;

(2) *it can handle concrete examples discussed in the philosophical literature*;

(3) *it goes beyond particular examples to provide a systematic and general view of the topic*;

⋆ Published in *Logic Across the University: Foundations and Applications* (College Publications), eds. Johan van Benthem and Fenrong Liu, 23-31.

(4) *it leads to philosophically-relevant discoveries that would be difficult to make by non-formal methods alone.*

Points (1) and (2) address the worry that by formalizing we may "change the subject." Points (3) and (4) address the question, "What do we get out of this?" Of course, there are other features that contribute to the value of an application of formal methods, as well as features that detract from it. Here I will explain the ways in which I think EC&ELI satisfies the four desiderata above.

EC&ELI proposes formalizations of the RA theories of Lewis [1996] and Heller [1989, 1999]; one way of developing the RA theory of Dretske [1981], based on Heller; and basic versions of the "subjunctivist" *tracking* theory of Nozick [1981] and *safety* theory of Sosa [1999]. To formalize the theories of Lewis and (one way of developing) Dretske, I introduce *RA models*, defined below. Traditional semantics for epistemic logic represent knowledge by the elimination of possibilities, but without any explicit distinction between those *relevant* possibilities that an agent must rule out in order to have knowledge vs. those remote, far-fetched or otherwise irrelevant possibilities that can be properly ignored. RA models make this distinction explicit by adding *relevance orderings*.

Definition 1. A *relevant alternatives* (RA) *model* is a tuple $\mathcal{M} = \langle W, \rightarrowtriangle, \preceq, V\rangle$ where $\langle W, \rightarrowtriangle, V\rangle$ is an ordinary relational model, with $\rightarrowtriangle$ at least reflexive, and $\preceq$ assigns to each $w \in W$ a *preorder* (reflexive and transitive) $\preceq_w$ on some $W_w \subseteq W$, for which w is a minimal element. Let $u \prec_w v$ iff $u \preceq_w v$ and not $v \preceq_w u$, and define $\mathrm{Min}_{\preceq_w}(S) = \{v \in S \cap W_w \mid \text{there is no } u \in S \text{ such that } u \prec_w v\}$.

Intuitively, $w \rightarrowtriangle v$ means that v is an *uneliminated* epistemic alternative for the agent in w; $u \prec_w v$ means that u is a *more relevant alternative* at w than v is; and $\mathrm{Min}_{\preceq_w}(S)$ is the set of *most relevant alternatives* in S at w. For simplicity, assume that each $\preceq_w$ is well-founded: $S \cap W_w \neq \emptyset$ implies $\mathrm{Min}_{\preceq_w}(S) \neq \emptyset$. (The main results also hold with more general definitions without well-foundedness.)

Now we can define three semantics for the basic epistemic language with formulas $\varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid K\varphi$: C-semantics for Cartesian, D-semantics for (one way of developing) Dretske [1981], and L-semantics for Lewis [1996].

Definition 2. Given a well-founded RA model $\mathcal{M} = \langle W, \rightarrowtriangle, \preceq, V\rangle$ with $w \in W$ and a formula φ, define $\mathcal{M}, w \vDash_x \varphi$ and $[\![\varphi]\!]_x^{\mathcal{M}} = \{v \in W \mid \mathcal{M}, v \vDash_x \varphi\}$ as follows (with propositional cases as usual):

$$\begin{aligned}
\mathcal{M}, w \vDash_c K\varphi \quad &\text{iff} \quad \forall v \in \overline{[\![\varphi]\!]_c^{\mathcal{M}}} : w \not\rightarrowtriangle v;\\
\mathcal{M}, w \vDash_d K\varphi \quad &\text{iff} \quad \forall v \in \mathrm{Min}_{\preceq_w}\big(\overline{[\![\varphi]\!]_d^{\mathcal{M}}}\big) : w \not\rightarrowtriangle v;\\
\mathcal{M}, w \vDash_l K\varphi \quad &\text{iff} \quad \forall v \in \mathrm{Min}_{\preceq_w}(W) \cap \overline{[\![\varphi]\!]_l^{\mathcal{M}}} : w \not\rightarrowtriangle v,
\end{aligned}$$

where $\overline{P} = \{v \in W \mid v \notin P\}$. In C-semantics, for an agent to know φ in world w, *all* of the $\neg\varphi$-possibilities must be eliminated by the agent in w. In D-semantics, for any φ there is a set $\mathrm{Min}_{\preceq_w}\big(\overline{[\![\varphi]\!]_d^{\mathcal{M}}}\big)$ of *most relevant* (*at* w) $\neg\varphi$-possibilities that the agent must eliminate in order to know φ. Finally, in L-semantics, there is a set of relevant possibilities, $\mathrm{Min}_{\preceq_w}(W)$, such that for any φ, in order to know φ the agent must eliminate the $\neg\varphi$-possibilities *within that set*.

Observe that for D-semantics, the whole relevance ordering $\preceq_w$ matters. As Heller puts it, there are "worlds surrounding the actual world ordered according to how realistic they are, so that those worlds that are more realistic are closer to the actual world than the less realistic ones" [1989, 25] with "those that are too far away from the actual world being irrelevant" [1999, 199], where how far is "too far" depends on the φ in question. By contrast, L-semantics does not use the ordering of more or less relevant worlds beyond the set $\mathrm{Min}_{\preceq_w}(W)$, taken to represent Lewis's single set of relevant alternatives in the current context (for a formal treatment of the dynamics of *context change*, see Holliday 2012a).

To formalize the RA theory of Heller [1989, 1999], the basic *tracking* theory of Nozick [1981], and the basic *safety* theory of Sosa [1999], I introduce CB models:

Definition 3. A *counterfactual belief* (CB) *model* is a tuple $\mathcal{M} = \langle W, D, \leqslant, V\rangle$ where W, $\leqslant$, and V are defined like W, $\preceq$, and V in Definition 1, and D is a serial binary relation on W.

D is a doxastic accessibility relation, so that wDv means that everything the agent believes in w is true in v. To make the truth clauses for K clearer, I add a belief operator B for the D relation; but the main results will be stated for the epistemic language without B. The preorders $\leqslant_w$ can be thought of either as relevance orderings, as before, or as *similarity orderings* as in Lewis's [1973] semantics for counterfactuals. With this setup, we can define three more semantics, formalizing three "subjunctivist" views of knowledge: H-semantics for Heller [1999], N-semantics for Nozick [1981], and S-semantics for Sosa [1999]. There are a number of qualifications to be made here, but I will mention only three, referring to the full paper for the rest. First, in EC&ELI, I only treat the basic versions of the tracking and safety theory (for the versions with "methods" and "bases" of belief, see Holliday 2012b, §2.D). Second, Heller and Sosa only propose necessary conditions for knowledge, so they may wish to take K here to represent sensitive belief or safe belief, rather than knowledge. Third, I interpret the sensitivity, adherence, and safety conditions along Lewisian counterfactual lines, which is common in the literature, but not the only conceivable interpretation.

Definition 4. Given a well-founded CB model $\mathcal{M} = \langle W, D, \leqslant, V\rangle$ with $w \in W$ and formula φ, define $\mathcal{M}, w \vDash_x \varphi$ and $[\![\varphi]\!]^{\mathcal{M}}_x = \{v \in W \mid \mathcal{M}, v \vDash_x \varphi\}$ as follows:

$\mathcal{M}, w \vDash_x B\varphi$ iff $\forall v \in W$: if wDv then $\mathcal{M}, v \vDash_x \varphi$;

$\mathcal{M}, w \vDash_h K\varphi$ iff $\mathcal{M}, w \vDash_h B\varphi$ and
(sensitivity) $\forall v \in \mathrm{Min}_{\leqslant_w}\big(\overline{[\![\varphi]\!]^{\mathcal{M}}_h}\big)$: $\mathcal{M}, v \nvDash_h B\varphi$;

$\mathcal{M}, w \vDash_n K\varphi$ iff $\mathcal{M}, w \vDash_n B\varphi$ and
(sensitivity) $\forall v \in \mathrm{Min}_{\leqslant_w}\big(\overline{[\![\varphi]\!]^{\mathcal{M}}_n}\big)$: $\mathcal{M}, v \nvDash_n B\varphi$, and
(adherence) $\forall v \in \mathrm{Min}_{\leqslant_w}\big([\![\varphi]\!]^{\mathcal{M}}_n\big)$: $\mathcal{M}, v \vDash_n B\varphi$;

$\mathcal{M}, w \vDash_s K\varphi$ iff $\mathcal{M}, w \vDash_s B\varphi$ and
(safety) $\forall v \in \mathrm{Min}_{\leqslant_w}\big([\![B\varphi]\!]^{\mathcal{M}}_s\big)$: $\mathcal{M}, v \vDash_s \varphi$.

In H-semantics, an agent knows φ iff she believes φ in the actual world but not in any of the "closest" (most similar or relevant) $\neg\varphi$-worlds. This condition implies both that φ is true in the actual world and, according to Lewis's semantics for counterfactuals, that *if φ were false, the agent would not believe φ*; in this sense, her belief is "sensitive" to the truth of φ. In N-semantics, there is an additional requirement for the agent to know φ, namely that she believes φ in all of the closest φ-worlds; in this sense, her belief is "adherent" to the truth of φ.[1] Finally, in S-semantics, an agent knows φ iff she believes φ and in the closest worlds where she believes φ, φ is true; in this sense, her belief in φ is "safe." As explained in EC&ELI, there are structural similarities between D-semantics and H/N-semantics, and between L-semantics and S-semantics.

By drawing RA and CB models as in Figs. 1 - 2,[2] one can represent many concrete examples from the epistemological literature, ranging from examples meant to challenge epistemic closure to examples meant to challenge the necessity or sufficiency of the subjunctivist conditions for knowledge. Doing so helps make clear what exactly one must assume for the challenge to be successful. Indeed, it may be helpful if epistemologists were to draw such diagrams when presenting putative counterexamples.

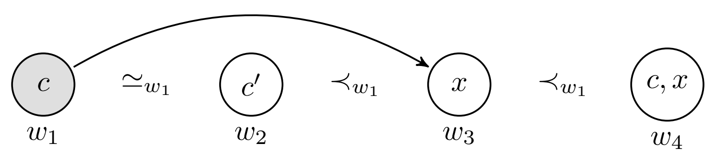

Fig. 1. RA model for Example 1 in EC&ELI (partially drawn, reflexive loops omitted)

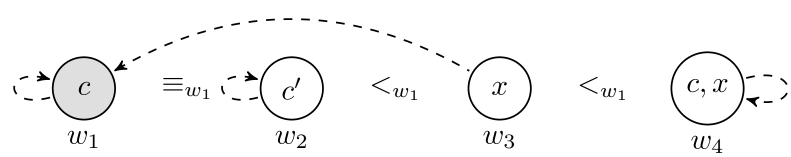

Fig. 2. CB model for Example 1 in EC&ELI (partially drawn)

[1] *Tracking* is the conjunction of sensitivity and adherence. Nozick used 'variation' for what I call 'sensitivity' and used 'sensitivity' to cover both variation and adherence; but the narrower use of 'sensitivity' is now standard. For discussion of different options for formalizing adherence, see Observation 4.5 of EC&ELI.

[2] The solid arrow in the RA model represents the $\rightarrow$ relation, and the dashed arrows in the CB model represents the D relation; $w_1 \simeq_{w_1} w_2$ indicates that $w_1 \preceq_{w_1} w_2$ and $w_2 \preceq_{w_1} w_1$, and $w_1 \equiv_{w_1} w_2$ indicates that $w_1 \leqslant_{w_1} w_2$ and $w_2 \leqslant_{w_1} w_1$; and just as $\prec_{w_1}$ is the strict part of $\preceq_{w_1}$, $<_{w_1}$ is the strict part of $\leqslant_{w_1}$.

C/D/L/H/N/S-semantics formalize views of knowledge that have been discussed extensively in epistemology. Having formalized these views, we can move beyond particular examples to answer general questions. Consider, for instance, the question of epistemic closure. Given sequences of formulas $\varphi_1, \dots, \varphi_n$ and $\psi_1, \dots, \psi_m$, and a propositional conjunction φ_0, let us write

$$\chi_{n,m} := \varphi_0 \wedge K\varphi_1 \wedge \cdots \wedge K\varphi_n \to K\psi_1 \vee \cdots \vee K\psi_m.$$

Call such a $\chi_{n,m}$ a *closure principle*. It states that if the agent knows each of φ_1 through φ_n (and the world satisfies a non-epistemic φ_0), then the agent knows at least one of ψ_1 through ψ_m. Our question is: which closure principles are *valid*?

Theorem 1 below provides the answer. It is an example of a discovery that we can take back to epistemology, illuminating the closure properties of knowledge according to standard theories, which would be difficult to make without formal investigation. Surprisingly, despite the differences between the RA, tracking, and safety theories of knowledge as formalized by D/H/N/S-semantics, Theorem 1 provides a unifying perspective: *the valid epistemic closure principles are essentially the same for these different theories*. The only twist is with D-semantics over *total* RA models, i.e., RA models in which for all $w \in W$, $\preceq_w$ is a total preorder on W_w, meaning that all possibilities are comparable in relevance. For comparison, I also include C/L-semantics, which fully support closure in the sense that if $\varphi_1 \wedge \cdots \wedge \varphi_n \to \psi$ is valid, then so is $K\varphi_1 \wedge \cdots \wedge K\varphi_n \to K\psi$.

The statement of Theorem 1 refers to a "T-unpacked" closure principle, which I will not fully define here (see EC&ELI, §5.2.1). For the first reading of the theorem, think only of "flat" closure principles $\chi_{n,m}$ without nesting of the K operator, which are T-unpacked if φ_0 contains $\varphi_1 \wedge \cdots \wedge \varphi_n$ as a conjunct. A key fact is that any formula in the basic epistemic language is equivalent to a conjunction of T-unpacked closure principles (see §5.2.1), so if we can decide the validity of such principles, then we can decide the validity of any formula.

Theorem 1 (Closure Theorem). *Let*

$$\chi_{n,m} := \varphi_0 \wedge K\varphi_1 \wedge \cdots \wedge K\varphi_n \to K\psi_1 \vee \cdots \vee K\psi_m$$

be a T-unpacked closure principle.

1. $\chi_{n,m}$ *is C/L-valid over* relevant alternatives *models iff*
   - *(a)* $\varphi_0 \to \bot$ *is valid or*
   - *(b) for some* $\psi \in \{\psi_1, \dots, \psi_m\}$,

     $$\varphi_1 \wedge \cdots \wedge \varphi_n \to \psi \text{ is valid;}$$

2. $\chi_{n,m}$ *is D-valid over* total relevant alternatives *models iff (a) or*
   - *(c) for some* $\Phi \subseteq \{\varphi_1, \dots, \varphi_n\}$ *and nonempty* $\Psi \subseteq \{\psi_1, \dots, \psi_m\}$,

     $$\bigwedge_{\varphi \in \Phi} \varphi \leftrightarrow \bigwedge_{\psi \in \Psi} \psi \text{ is valid;}$$

3. $\chi_{n,m}$ *is D-valid over* all relevant alternatives *models iff (a) or*
   (d) *for some* $\Phi \subseteq \{\varphi_1, \dots, \varphi_n\}$ *and* $\psi \in \{\psi_1, \dots, \psi_m\}$,

$$\bigwedge_{\varphi \in \Phi} \varphi \leftrightarrow \psi \textit{ is valid.}$$

4. $\chi_{n,m}$ *is H/N/S-valid over* counterfactual belief *models if (a) or (d); and a* flat $\chi_{n,m}$ *is H/N/S-valid over such models only if (a) or (d).*[3]

I will discuss some logical aspects of Theorem 1 before its epistemological upshot. Using Theorem 1, we can reduce the validity of any closure principle to the validity of finitely many formulas of lesser modal depth. With a "modal decomposition" result of this form, van Benthem [2010] proves completeness of the weakest normal modal logic **K** with respect to relational models (and notes the analogous result for the weakest monotonic modal logic **EM** with respect to monotonic neighborhood models). Like van Benthem's proof, the proof of Theorem 1 in the 'only if' directions assumes that the formulas of lesser modal depth are not valid, from which we infer the existence of models to "glue together" into a countermodel for $\chi_{n,m}$. But from there the proof of Theorem 1 requires new techniques. First, since we are dealing with models in which $K\varphi \to \varphi$ is valid, we must use the new idea of T-unpacking. Second, since we are dealing with a hybrid of relational and *ordering* semantics, we cannot simply glue all of the relevant models together at once, as in the basic relational case; instead, we must put them in the right order, which can be done inductively (see §5.2.2).

Using Theorem 1 we obtain the axiomatizations in Corollary 1 below, stated for the axioms and rules in Table 1. **E** is the weakest of the *classical* modal systems with PL, MP, and RE. $\mathbf{ES}_1 \dots \mathbf{S}_n$ is the extension of **E** with every instance of schemas $\mathrm{S}_1 \dots \mathrm{S}_n$. The logic **ECNTX**, which I dub *the logic of ranked relevant alternatives*, appears not to have been previously identified in the literature.

Corollary 1 (Soundness and Completeness).

1. ***EMCNT*** *(equivalently,* ***ET*** *plus RK, a.k.a.* ***KT****) is sound and complete for C/L-semantics over RA models.*
2. ***ECNTX*** *(equivalently,* ***ET*** *plus RAT) is sound and complete for D-semantics over* total *RA models.*
3. ***ECNT*** *(equivalently,* ***ET*** *plus RA) is sound and complete for D-semantics over RA models.*
4. ***ECNT*** *is sound (with respect to the full epistemic language) and complete (with respect to the flat fragment) for H/N/S-semantics over CB models (see §8.2 on higher-order knowledge).*[4]

[3] See §8.2 on higher-order knowledge for the subjunctivist H/N/S-semantics.

[4] Corollary 1.4 gives the answer, for the flat fragment of the epistemic language (without nesting of K), to the question posed by van Benthem [2010, 153] of what is the epistemic logic of Nozick's [1981] notion of knowledge.

PL. all tautologies	MP. $\dfrac{\varphi \to \psi \quad \varphi}{\psi}$
T. $K\varphi \to \varphi$ N. $K\top$	RE. $\dfrac{\varphi \leftrightarrow \psi}{K\varphi \leftrightarrow K\psi}$
M. $K(\varphi \wedge \psi) \to K\varphi \wedge K\psi$	RK. $\dfrac{\varphi_1 \wedge \cdots \wedge \varphi_n \to \psi}{K\varphi_1 \wedge \cdots \wedge K\varphi_n \to K\psi}$
X. $K(\varphi \wedge \psi) \to K\varphi \vee K\psi$	RAT. $\dfrac{\varphi_1 \wedge \cdots \wedge \varphi_n \leftrightarrow \psi_1 \wedge \cdots \wedge \psi_m}{K\varphi_1 \wedge \cdots \wedge K\varphi_n \to K\psi_1 \vee \cdots \vee K\psi_m}$
C. $K\varphi \wedge K\psi \to K(\varphi \wedge \psi)$	RA. $\dfrac{\varphi_1 \wedge \cdots \wedge \varphi_n \leftrightarrow \psi}{K\varphi_1 \wedge \cdots \wedge K\varphi_n \to K\psi}$

Table 1. axiom schemas and rules ($n \geq 0$, $m \geq 1$)

The epistemological upshot of Theorem 1 is that according to the family of basic RA, tracking, and safety theories of knowledge formalized by D/H/N/S-semantics, an ideal logician is guaranteed to know a logical consequence ψ of what she knows if and only if ψ *is equivalent to a conjunction of propositions she already knows*! By applying this test, one can check that not only is the strong closure principle $K\varphi \wedge K(\varphi \to \psi) \to K\psi$ invalid—which is arguably desirable—but so are very weak and uncontroversial closure principles, such as $K(\varphi \wedge \psi) \to K\varphi$ and $K\varphi \to K(\varphi \vee \psi)$, and even the likes of $K(\varphi \wedge \psi) \to K(\varphi \vee \psi)$ and $K\varphi \wedge K\psi \to K(\varphi \vee \psi)$, while more controversial principles like $K\varphi \wedge K\psi \to K(\varphi \wedge \psi)$ turn out valid. There is much more to be said about Theorem 1 and the ramifications of closure failures for higher-order knowledge (see EC&ELI, §8). But suffice it to say that I take Theorem 1 to be a serious negative result for the RA and subjunctivist theories in question,[5] which motivates the search for new and improved pictures of knowledge in the sequel to EC&ELI.[6]

The last point brings me back to the starting list of virtues in formalization, in order to add one more. An application of formal methods in philosophy also gains value if:

(5) *it leads to the development of new theories that solve philosophical problems.*

[5] There are other problems for C- and L-semantics having to do with skepticism and "vacuous knowledge." See EC&ELI, §3 and Holliday 2013c.

[6] However, we can also take Theorem 1 to be a neutral result about other desirable epistemic properties, viz., the properties of having ruled out the relevant alternatives to a proposition, of having a belief that tracks the truth of a proposition, or of having a safe belief in a proposition. For example, replace the K by a $\Box$, reading $\Box\varphi$ as "the agent has ruled out the most relevant alternatives to φ," and the newly identified logic **ECNTX**, *the logic of ranked relevant alternatives*, is of independent interest.

To develop such a theory in the case of knowledge is the goal of further work (Holliday 2013c,d, 2012b), for which EC&ELI lays the foundation.

There are many claims in EC&ELI about specific theories and results, summarized in the introduction and conclusion of the paper, but here I will end with some general methodological claims that I think EC&ELI supports:

- While "formal epistemology" often deals with different issues than "traditional" or "mainstream" epistemology deals, EC&ELI shows that a formal approach—in this case, a model-theoretic logical approach—can contribute to our understanding of central issues in traditional epistemology. (For more on epistemic logic and epistemology, see Holliday 2013a.)
- For modal logicians, epistemology represents an area of sophisticated theorizing in which modal-logical tools can help to clarify and systematize parts of the philosophical landscape.
- The attempt to apply modal-logical tools to epistemology also benefits modal logic by broadening its scope, bringing interesting new structures and systems under its purview.
- Some of these epistemological investigations may be like NASA space missions, pursued for their intrinsic interest but—as a bonus—also leading to new "technology" usable in other areas. In EC&ELI, proving the main results involved an alternative *modal decomposition* approach to modal completeness theorems, developed in novel ways for new logics. In future work, I plan to show the wider applicability of this approach to other modal logics.

Bibliography

J. v. Benthem. *Modal Logic for Open Minds*. CSLI Publications, Stanford, CA, 2010.

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

M. Heller. Relevant Alternatives. *Philosophical Studies*, 55(1):23–40, 1989.

M. Heller. Relevant Alternatives and Closure. *Australasian Journal of Philosophy*, 77(2):196–208, 1999.

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 Logic and Epistemology. In S. O. Hansson and V. F. Hendricks, editors, *Handbook of Formal Philosophy*. Springer, Dordrecht, 2013a. Forthcoming.

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

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

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

J. L. Kvanvig. Closure Principles. *Philosophy Compass*, 1(3):256–267, 2006.

D. Lewis. *Counterfactuals*. Basil Blackwell, Oxford, 1973.

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

R. Nozick. *Philosophical Explanations*. Harvard University Press, Cambridge, MA, 1981.

E. Sosa. How to Defeat Opposition to Moore. *Noûs*, 33(13):141–153, 1999.