Title
The Orthologic of Epistemic Modals

Authors
Holliday, Wesley Halcrow
Mandelkern, Matthew

Publication Date
2023-05-23

The Orthologic of Epistemic Modals

Wesley H. Holliday† and Matthew Mandelkern‡

† University of California, Berkeley (wesholliday@berkeley.edu)
‡ New York University (mandelkern@nyu.edu)

May 23, 2023

Abstract

Epistemic modals have peculiar logical features that are challenging to account for in a broadly classical framework. For instance, while a sentence of the form $p \wedge \Diamond\neg p$ ('p, but it might be that not p') appears to be a contradiction, $\Diamond\neg p$ does not entail $\neg p$, which would follow in classical logic. Likewise, the classical laws of distributivity and disjunctive syllogism fail for epistemic modals. Existing attempts to account for these facts generally either under- or over-correct. Some theories predict that $p \wedge \Diamond\neg p$, a so-called *epistemic contradiction*, is a contradiction only in an etiolated sense, under a notion of entailment that does not always allow us to replace $p \wedge \Diamond\neg p$ with a contradiction; these theories underpredict the infelicity of embedded epistemic contradictions. Other theories savage classical logic, eliminating not just rules that intuitively fail, like distributivity and disjunctive syllogism, but also rules like non-contradiction, excluded middle, De Morgan's laws, and disjunction introduction, which intuitively remain valid for epistemic modals. In this paper, we aim for a middle ground, developing a semantics and logic for epistemic modals that makes epistemic contradictions genuine contradictions and that invalidates distributivity and disjunctive syllogism but that otherwise preserves classical laws that intuitively remain valid. We start with an *algebraic semantics*, based on ortholattices instead of Boolean algebras, and then propose a *possibility semantics*, based on partial possibilities related by compatibility. Both semantics yield the same consequence relation, which we axiomatize. We then show how to lift an arbitrary possible worlds model for a non-modal language to a possibility model for a language with epistemic modals. Further, we show that we can use this construction to lift standard possible worlds treatments of probabilities and conditionals into possibility semantics. The goal throughout is to retain what is desirable about classical logic while accounting for the non-classicality of epistemic vocabulary.

Keywords: epistemic modals, negation, connectives, probability, conditionals, ortholattices, algebraic semantics, possibility semantics

1 Introduction

Exploration of epistemic modals in the last decades has shown that they do not fit easily into the framework of classical modal logic. Following Yalcin (2007), we can characterize the problem as follows. There is strong evidence that sentences of the form $p \wedge \Diamond\neg p$ or $\neg p \wedge \Diamond p$ (as well as variants that reverse the order) are contradictory: not only are they unassertable, but also they embed like contradictions across the board. To get a sense for the evidence, note the infelicity of the following:

(1) a. #The butler is the murderer, but he might not be.
b. #Suppose that the butler is the murderer, but he might not be.
c. #Everyone who is tall might not be tall.

In light of these data, it is natural to consider a logic that makes sentences of the form $p \wedge \Diamond\neg p$ contradictory. The problem is that adding $p \wedge \Diamond\neg p \vDash \bot$ as an entailment to any modal logic that extends classical logic immediately yields the entailment $\Diamond\neg p \vDash \neg p$. But $\Diamond\neg p$ plainly does not entail $\neg p$. 'The butler might not be the murderer' does not entail that the butler *is not* the murderer.

A profusion of responses to this problem have been developed in recent years. These responses either reject the assumption that $p \wedge \Diamond\neg p$ is truly a contradiction or hold that it is a contradiction only under a non-classical notion of $\vDash$, one that blocks the reasoning above. For a prominent example, consider the dynamic system of Groenendijk et al. (1996), wherein $p \wedge \Diamond\neg p$ is a contradiction,[1] but conjunction and entailment are interpreted non-classically, so that we cannot conclude that $\Diamond\neg p \vDash \neg p$ (Groenendijk et al. 1996). The non-classicality of this approach, however, goes very deep: for instance, the classical laws of non-contradiction and excluded middle also fail in the dynamic system, but we know of no evidence for their failure from modal language. Our goal in this paper is to develop a more minimalist non-classical approach to epistemic modals that validates $p \wedge \Diamond\neg p \vDash \bot$ without validating $\Diamond\neg p \vDash \neg p$. In particular, we block the inference from $p \wedge \Diamond\neg p \vDash \bot$ to $\Diamond\neg p \vDash \neg p$ by treating negation, algebraically speaking, as an *orthocomplementation* (the complement operation characteristic of ortholattices) but not necessarily a *pseudocomplementation* (the complement operation of Heyting and Boolean algebras, obeying the law that $a \wedge b = 0$ implies $b \leq \neg a$). However, we aim to depart from classical logic in a minimal way, by invalidating only those classical inference patterns that have intuitive counterexamples involving epistemic modals. We note in particular that those intuitive counterexamples always involve combinations of sentences with different levels of epistemic iteration: for example, when p itself is modal-free, $p \wedge \Diamond\neg p$ conjoins a sentence p with no epistemic modality and a sentence $\Diamond\neg p$ with one level of epistemic modality. Hence we develop a system that is fully classical when reasoning with sentences of the same "epistemic level." On this picture, classical reasoning is dangerous only when it crosses different epistemic levels.

We start in § 2 by drawing out intuitive desiderata for a logic of epistemic modals. We then characterize an *epistemic orthologic* that meets those desiderata, first algebraically in § 3 and then using a possibility semantics in § 4. Both semantics yield the same consequence relation. We provide a sound and complete axiomatization for this epistemic orthologic, following the call in Holliday and Icard 2018 to axiomatize proposed consequence relations in natural language semantics. In § 5, we show how to lift an arbitrary possible worlds model for a non-modal language to a possibility model for a language with epistemic modals, which yields an implementation of our possibility semantics built on the more familiar primitives of possible worlds semantics. Further, we show that this construction yields a natural way to simultaneously lift standard possible worlds treatments of probability and conditionals into possibility semantics, thus providing a first sketch of an attractive treatment of probabilities and conditionals in our framework. Finally, we compare our approach to existing approaches in § 6 and conclude in § 7.

An associated online repository (https://github.com/wesholliday/ortho-modals) contains a Jupyter notebook with code to use our logic with the Natural Language Toolkit's interface to the Prover9 theorem prover and Mace4 model builder, as well as a Jupyter notebook with code for using the possibility semantics of § 4.

[1] Provided p is free of modals or conditionals.

2 Desiderata

We will begin by bringing out the key desiderata for the logic of epistemic modals. In particular, our goal in this section and the next is to identify properties of an entailment relation appropriate to a language with epistemic modality. As a methodological preliminary, let us say a bit more about the kind of entailment relation we aim to characterize. Our target consequence relation is a classical one in the sense that it aims to capture universal preservation of truth. This has two upshots worth noting. First, φ entails ψ iff φ is semantically equivalent to $\varphi \wedge \psi$, where two formulas φ, ψ are semantically equivalent iff for any formulas χ and ρ, $\chi[\varphi/\rho]$ is true iff $\chi[\psi/\rho]$ is true (where $\chi[\varphi/\rho]$ is the sentence that results from uniformly substituting φ for ρ in χ; we will ignore potential issues about hyperintensionality here). Second, if φ entails ψ, then the probability of ψ must be at least that of φ.

We highlight these two upshots of the classical consequence relation because they yield two ways to empirically test our target notion of entailment. For a brief illustration, consider the inference from φ to $\Box\varphi$, which comes out valid in many logics for epistemic modals. We will invalidate this inference, and indeed, it is invalid in a logic that aims to capture inferences that universally preserve truth. First note that if φ entails ψ in the sense above, then the probability of ψ will always be at least as great as the probability of φ (since on any natural notion of probability, the probability of semantically equivalent propositions will be equal, and the probability of a conjunct will always be at least as great as the probability of a conjunction). But then if φ entailed $\Box\varphi$, the probability of $\Box\varphi$ would always have to be at least as great as that of φ. But this is wrong. Consider a fair coin that was just flipped; suppose we do not know the outcome of the flip. Plausibly, the probability that it landed heads is around .5, but the probability that it *must* have landed heads is much lower, around 0. Second, since the present notion of entailment contraposes (given classical assumptions about De Morgan's laws and negation that we will not question here), if $\varphi \vDash \Box\varphi$, then $\neg\Box\varphi \vDash \neg\varphi$; but that is plainly implausible, since $\neg\Box\varphi$ is equivalent to $\Diamond\neg\varphi$, but clearly $\Diamond\neg\varphi$ does not entail $\neg\varphi$ (from 'It might not be raining' it does not follow that it is not raining). Finally, given that most accept that $\Box\varphi$ entails φ, if φ also entailed $\Box\varphi$ then they would be logically equivalent. But then they would be everywhere substitutable for one another, salva veritate. However, they are clearly not: $\neg\varphi$ and $\neg\Box\varphi$ are not equivalent (compare 'it's not raining' to 'it's not the case that it *must* be raining').

There are other kinds of entailment relations that we can also characterize: for instance, prominently, we might develop a logic characterizing the inferences that preserve rational acceptance in all cases (Yalcin 2007, Bledin 2014, 2015, Santorio 2021, Norlin 2020, 2022). These logics are well worth studying but are not our central target here. In the conclusion, we briefly discuss how to characterize a logic of acceptance in our framework. In such a logic—unlike ours—the inference from φ to $\Box\varphi$ might reasonably be seen to be valid. However, we do not see any sense in which our project (characterizing a logic of truth preservation) and the project of characterizing a logic of acceptance are in conflict. There is a purely terminological question about which of these deserves the name of 'logic', a question whose interest we do not see; we are happy to use 'logic' in the traditional way for the set of truth-preserving inferences, but if others prefer different usages, feel free to substitute your preferred terminology. There could also be a substantive question if someone thought that we should *only* characterize the logic of acceptance and that the logic of truth-preservation is uninteresting or irrelevant. But such a position has little merit, and (unsurprisingly) has not been defended in the literature to our knowledge. For even if we had an adequate characterization of the logic of acceptance, there would remain the further question of which inferences preserve probability and of which sentences are always substitutable salva veritate. These questions are not answered by a logic of acceptance (since we might have, e.g., that φ entails $\Box\varphi$ on that notion, but the inference does not preserve probability, nor are

φ and $\Box\varphi$ substitutable). So everyone interested in core semantic notions needs to characterize the logic of truth preservation. It is natural to think that a logic of rational acceptance will then supervene on this logic, together with a notion of rational acceptance (and that is indeed how such a logic is defined in, e.g., Yalcin 2007). But even if you reject this supervenience claim, you still need a logic of truth preservation, since it is not plausible that the logic of truth preservation supervenes on the logic of rational acceptance.

We take these points to be uncontroversial, but questions about our target notion of entailment have frequently been raised in conversation and are particularly pertinent to the literature on epistemic modals (where it is often claimed that φ entails $\Box\varphi$—which, again, we are happy to accept in one sense, but not in our target sense of entailment). Let us now turn to our substantive desiderata for our target logic.

2.1 Epistemic contradictions

First, we give more evidence that sentences of the form $p\wedge\Diamond\neg p$ and $\neg p\wedge\Diamond p$ are indeed contradictory, drawing on observations in Groenendijk et al. 1996, Aloni 2000, Yalcin 2007, 2015, and Mandelkern 2019. Yalcin calls sentences of this form *epistemic contradictions*. Mandelkern (2019) calls epistemic contradictions and variants that reverse their order (that is, sentences of the form $\Diamond\neg p\wedge p$ or $\Diamond p\wedge\neg p$) *Wittgenstein sentences*. Although this is somewhat controversial (and runs counter to the prominent dynamic approach mentioned above), we think that order does not matter to our central points, so we intend the claims we make here to apply to all Wittgenstein sentences. However, for brevity, we often use epistemic contradictions as a stand-in for all Wittgenstein sentences.

To start, note that Wittgenstein sentences are generally unassertable. It is hard to think of circumstances in which any of the following can be asserted:

(2) a. #It's raining, but it might not be.
b. #It might be raining, but it isn't.
c. #The cat might be under the bed, but she is on the couch.
d. #I'll make lasagna, but I might not make lasagna.
e. #Sue isn't the winner but she might be.

Unassertability, however, could plausibly be due to either pragmatic or semantic factors. Indeed, these sentences are superficially similar to Moore sentences (as Wittgenstein (1953) observed), and the latter are standardly explained on a pragmatic basis: you cannot assert, say, 'It's raining, but I don't know that', because you cannot know this, on pain of contradiction, and you need to know what you assert.[2]

To test whether the unassertability of Wittgenstein sentences is pragmatic or semantic, we can look at how they embed in various environments, comparing their behavior with embedded Moore sentences on the one hand and embedded classical contradictions on the other, building on the structure of arguments of Yalcin 2007. So consider the three-way comparisons below, with first the Moore sentence, then the Wittgenstein sentence, and finally the contradiction:

(3) a. Suppose Sue is the winner but I don't know it.
b. #Suppose Sue is the winner but she might not be.
c. #Suppose Sue is the winner and isn't the winner.

(4) a. If Sue is the winner but I don't know it, I will chastise her for not telling me.

[2] Cf. Hintikka 1962, § 4.11. For a study of Moore sentences in the context of epistemic logic and dynamic epistemic logic, see Holliday and Icard 2010.

b. #If Sue is the winner but she might not be, I will chastise her for not telling me.

c. #If Sue is the winner and isn't the winner, I will chastise her for not telling me.

(5) a. It could be that Sue is the winner but I don't know it.

b. #It could be that Sue is the winner and might not be.

c. #It could be that Sue is the winner and isn't the winner.

(6) a. Either Sue is the winner but I don't know it, or she isn't the winner but I don't know it.

b. #Either Sue is the winner but might not be, or she isn't the winner but might be.

c. #Either Sue is the winner and isn't the winner, or she isn't the winner and is the winner.

(7) a. The winner is, for all I know, not the winner.

b. #The winner might not be the winner.

c. #The winner isn't the winner.

(8) a. Someone is the winner and for all I know isn't the winner.

b. #Someone is the winner and might not be the winner.

c. #Someone is the winner and isn't the winner.

In all these cases, the first variant, embedding a Moore sentence, is felicitous, supporting the idea that the infelicity of Moore sentences is pragmatic. By contrast, both of the latter variants are infelicitous, suggesting that the infelicity of Wittgenstein sentences is not pragmatic, after all, but instead due to the fact that they are genuine contradictions.

This is not to say that Wittgenstein sentences and classical contradictions pattern in exactly the same ways. It is well known that classical contradictions are often judged by ordinary speakers to be interpretable ('Sue is the winner and isn't the winner' could be used to express that Sue is the winner in one sense but not in some other sense), and Wittgenstein sentences, embedded or not, can similarly be coerced (though perhaps with more difficulty) into slightly different communicative functions. Nonetheless, their behavior across the board is strikingly like that of contradictions. We conclude that they are indeed contradictions and that any residual differences between $p \wedge \neg p$ and $p \wedge \Diamond\neg p$ should be accounted for on the basis of their underlying compositional semantics.

This presents us with our first desideratum: a logic on which Wittgenstein sentences are contradictions. That is, we want $\varphi \wedge \Diamond\neg\varphi \vDash \bot$ (and likewise $\neg\varphi \wedge \Diamond\varphi \vDash \bot$, $\Diamond\varphi \wedge \neg\varphi \vDash \bot$, and $\Diamond\neg\varphi \wedge \varphi \vDash \bot$). In fact, we need more than this: we need a system in which $\varphi \wedge \Diamond\neg\varphi$ can always be replaced with $\varphi \wedge \neg\varphi$ *salva veritate*. Otherwise, we would not have an immediate account of the data above, even if $\varphi \wedge \Diamond\neg\varphi \vDash \bot$. This is worth mentioning since some systems, like domain semantics (Yalcin 2007), predict that $\varphi \wedge \Diamond\neg\varphi \vDash \bot$ but not that $\varphi \wedge \Diamond\neg\varphi$ can always be replaced with $\varphi \wedge \neg\varphi$ *salva veritate* and hence miss some of the data above (see § 6 for further discussion). Call a logic an *E-logic* (suggesting *epistemic*) if it satisfies these first two desiderata.

We have already seen why an E-logic is hard to obtain: given classical assumptions, treating Wittgenstein sentences as contradictions would make $\Diamond\neg p$ entail $\neg p$. So, in particular, we want a logic that is an E-logic but that does not yield the absurd conclusion that $\Diamond\neg p \vDash \neg p$. Lest this problem be treated as having to do essentially with the peculiarities of 'and', we should note that the problem can equally be formulated without involving conjunction at all. If we take the evidence above to show that p and $\Diamond\neg p$ are jointly inconsistent, then what we need is a logic on which $\{\Diamond\neg p, p\} \vDash \bot$. Classical reasoning would yield $\Diamond\neg p \vDash \neg p$. So we need a logic that blocks this reasoning. Of course, evidence that p and $\Diamond\neg p$ are jointly inconsistent is somewhat less direct than the evidence that $p \wedge \Diamond\neg p$ is itself inconsistent. After all, we cannot embed a pair of propositions under a unary sentential operator, as we did with the corresponding conjunction. But the sentences above already yield indirect evidence that it is not just the conjunction that is inconsistent, but also its conjuncts are jointly inconsistent, since the infelicity persists when the conjuncts are distributed across the restrictor and scope of quantifiers, as in (7) and (8). Similar evidence comes from pairs of attitude ascriptions, where again the infelicity persists when the conjuncts in question are distributed under distinct attitude predicates:

(9) a. #Liam believes that Sue is the winner. Liam also believes that Sue might not be the winner.
b. #Suppose that Sue is the winner. Suppose, further, that Sue might not be the winner.
c. #I hope Susie wins. I also hope she might not win.

Finally, another way to see the puzzle of epistemic contradictions, also emphasized by Yalcin (2007), is to approach it from intuitions about meaning and synonymy rather than intuitions about logic. It is very natural to think that $\Diamond p$ means something like 'For all we know, p is true'. This is roughly the meaning it has in the influential approach of Kratzer 1977, 1981, for instance. But if that is right, then $p \wedge \Diamond\neg p$ should have the meaning of a Moore sentence and should be embeddable just like Moore sentences are. But the examples above clearly show that this is not the case. So apparently $\Diamond p$ does not mean even roughly the same thing as 'For all we know, p is true'. But then what does it mean?

2.2 Distributivity

We have seen one reason to think that an adequate treatment of epistemic modals calls for revision to classical logic. A second reason, brought out in Mandelkern 2019, has to do with the distributive law, according to which $\varphi \wedge (\psi \vee \chi)$ is logically equivalent to $(\varphi \wedge \psi) \vee (\varphi \wedge \chi)$ for any sentences φ, ψ, χ. This is, of course, a law of classical logic. But while it is intuitively valid for modal-free fragments, it is not intuitively valid once epistemic modals are on the scene; for compare:

(10) a. Sue might be the winner and she might not be, and either she is the winner or she isn't.
b. #Sue might not be the winner and she is the winner, or else Sue might be the winner and she isn't the winner.

While (10-a) sounds simply like a long-winded avowal of ignorance, (10-b) is very strange. But if distributivity is valid, then (10-a) entails (10-b)! (In fact, given distributivity, (10-a) and (10-b) are logically *equivalent* provided that p entails $\Diamond p$, on which more shortly.) Such examples motivate us to look for a logic that invalidates distributivity for the modal fragment.

The ramifications of the failure of distributivity are plausibly very wide. For instance, these failures might also help explain a puzzling observation, which is based on discussion in Ninan 2018 and related discussion in Aloni 2000. Consider a fair lottery where at least one ticket will win, but not all will win. The winning ticket(s) have been drawn, but we do not know which tickets won. Then the following seems true:

(11) Every ticket might not be a winning ticket.

But we cannot conclude:

(12) Some winning ticket might not be a winning ticket.

The winning tickets, however, are among the tickets, and so if every ticket might not be a winning ticket, it seems that we should be able to conclude that some winning ticket might not be a winning ticket.

We can reconstruct this inference as follows. Suppose we have rigid designators for all the tickets: $t_1, \ldots, t_n$. Then (11) can be taken to be the conjunction $\Diamond\neg W(t_1) \wedge \cdots \wedge \Diamond\neg W(t_n)$, where W stands for 'is a winner'. We also know the disjunction $W(t_1) \vee \cdots \vee W(t_n)$. Putting these two together, we have

$$\big(\Diamond\neg W(t_1) \wedge \cdots \wedge \Diamond\neg W(t_n)\big) \wedge \big(W(t_1) \vee \cdots \vee W(t_n)\big);$$

in short, every ticket might not be a winner, but some ticket is a winner. Distributivity would allow us to infer

$$\big(W(t_1) \wedge \Diamond\neg W(t_1)\big) \vee \cdots \vee \big(W(t_n) \wedge \Diamond\neg W(t_n)\big),$$

that is, 'Some winning ticket might not be a winning ticket', as in (12). However, if distributivity is not valid, we cannot conclude (12) from (11). We will set aside quantification in this paper, but this example illustrates the importance of distributive reasoning—and identifying its failures—across a variety of examples involving epistemic modals.

2.3 Disjunctive syllogism

A closely related point is that disjunctive syllogism intuitively fails for epistemic modals (Klinedinst and Rothschild 2012, citing Yalcin). Disjunctive syllogism says that $\{p \vee q, \neg q\} \vDash p$. Hence the following inference, varying an example from Klinedinst and Rothschild, is valid if disjunctive syllogism is:

(13) a. Either the dog is inside or it must be outside.
b. It's not the case that the dog must be outside.
c. Therefore, the dog is inside.

This inference is intuitively invalid. For (13-a) feels a lot like the corresponding tautology:

(14) The dog is inside or outside.

And indeed, in any E-logic that validates De Morgan's laws (plus double negation elimination and duality for $\Diamond$ and $\Box$), (13-a) is a logical truth: $p \vee \Box\neg p$ is equivalent to $\neg(\neg p \wedge \Diamond p)$, which, if the logic is an E-logic, is equivalent to $\neg\bot$. Again, this seems intuitive: that is, sentences of the form $p \vee \Box\neg p$ do feel a lot like the corresponding tautologies $p \vee \neg p$, as in (15) (an intuition that our system will capture by predicting that $p \vee \Box\neg p$ is indeed a tautology):

(15) a. Either Fatima is home or else she must be somewhere else.
b. Either John must be the culprit or else it isn't him.

Thus, we know (13-a) just on the basis of the logic of epistemic modals. And (13-b) (which is equivalent, by duality, to 'the dog might be inside') can be true without the dog in fact being inside. (Given that (13-a) has the status of a logical truth, and given duality, this is just the point, again, that $\Diamond p$ does not entail p.) Hence disjunctive syllogism is not valid for a fragment including epistemic modals.

Given very weak assumptions, disjunctive syllogism follows from distributivity. So the intuitive failure of disjunctive syllogism gives us another reason to invalidate distributivity.

2.4 Orthomodularity

Another law of classical logic—and indeed of the weaker system of quantum logic (see, e.g., Chiara and Giuntini 2002)—that we have reason to invalidate is orthomodularity, according to which if $\varphi \vDash \psi$, then $\psi \vDash \varphi \vee (\neg\varphi \wedge \psi)$.[3] It is easy to see that this is classically valid (indeed, we have $\psi \vDash \varphi \vee (\neg\varphi \wedge \psi)$ in classical logic, though the assumption of $\varphi \vDash \psi$ is needed in quantum logic). But, building on observations in Dorr and Hawthorne 2013, we argue that it is invalid for epistemic modals.[4]

We assume, again, that $p \vDash \Diamond p$. By contraposition of $\vDash$ and duality, this follows from the assumption that 'must' is *factive*, i.e., that $\Box p \vDash p$, which is widely (though not universally) accepted; see von Fintel and Gillies 2010, 2021 for a variety of arguments for factivity, and see Footnote 5 for direct arguments that $p \vDash \Diamond p$.[5] Given that $p \vDash \Diamond p$, orthomodularity says that $\Diamond p \vDash p \vee (\neg p \wedge \Diamond p)$. But, if our logic is E, the right disjunct of $p \vee (\neg p \wedge \Diamond p)$ is contradictory, so the disjunction is equivalent to p; likewise, $(\neg p \wedge \Diamond p) \vee p$ is predicted to be equivalent to p. For example, consider (21):

(21) a. Either it isn't the butler but it might be, or else it's the butler.
b. It's the butler.

Intuitions here are somewhat unclear, because (21-a) sounds infelicitous—as we would expect a disjunction with a contradictory disjunct to sound. Nonetheless, as Dorr and Hawthorne (2013) observe, (21-a) feels intuitively equivalent to (21-b). However, orthomodularity predicts that 'It might be the butler' entails (21-a). Hence, given the contradictoriness of $\neg p \wedge \Diamond p$, orthomodularity would let us infer p from $\Diamond p$, which again is unacceptable.

Distributivity entails orthomodularity, so the intuitive failure of the latter gives us yet another reason to reject distributivity.

[3]That $\varphi \vDash \psi$ implies $\varphi \vee (\neg\varphi \wedge \psi) \vDash \psi$ holds by disjunction and conjunction elimination, in which case orthomodularity can be stated as: if $\varphi \vDash \psi$, then $\psi \dashv\vDash \varphi \vee (\neg\varphi \wedge \psi)$.

[4]Orthomodularity is a weakening of *modularity*, which says that if $\varphi \vDash \psi$ then $(\varphi \vee \chi) \wedge \psi \dashv\vDash \varphi \vee (\chi \wedge \psi)$. Hence related counterexamples also shows that modularity fails; for instance, as we argue presently, 'It's raining' entails 'It might be raining'; but 'It's raining or it's not raining, and it might be raining' is not plausibly equivalent to 'It's raining, or it's not raining and it might be raining'. See Remark 3.14 where related issues arise.

[5]Consider first proof by cases. The following seems valid:

(16) Either John is home or he is at work. So he might be at home or he might be at work.

Likewise, Moorean constructions like (17) are infelicitous, which is naturally explained if $p \vDash \Diamond p$:

(17) #It's raining, but I don't know whether it might be raining.

Retraction data display a similar pattern:

(18) The butler did it. . . .
a. #I'm not sure whether the butler might have done it, but what I said earlier is true.
b. #I only said that the butler *did it*, not that he *might have* done it.

The inference in question also seems to preserve truth under upward-entailing attitude predicates:

(19) John knows that Sue is at the party. So, John knows that Sue might be at the party.

Finally, the inference seems to be supported by Modus Ponens (which is generally accepted to be valid for conditionals with epistemic antecedents, even if, as McGee (1985) and following maintain, it is not valid for non-Boolean consequents). Suppose Mary says to Mark:

(20) If you might be sick, please let me know.

Now suppose Mark is in fact sick. In light of (20), it seems that he is therefore obligated to let Mary know that he is sick.

2.5 Conservativity

Our last desideratum is a kind of methodological conservatism. Where there is not a specific argument for the failure of a classically valid inference pattern, we should validate it, on the presumption that it is indeed valid. It is of course harder to argue for the validity of a schema than against it; arguing that a principle is invalid only requires a single convincing counterexample, whereas we cannot ever look at every instance of a schema by way of arguing for its validity. Still, we think that as a methodological principle it is sensible to proceed by minimally altering classical logic in light of counterexamples and examining the result. Thus, we will look for a minimal variation on classical (modal) logic that can meet the desiderata above—that is, which is an E-logic and invalidates distributivity and even orthomodularity for the modal fragment, but still validates these, along with all other classical laws, for the non-modal fragment. Across the modal fragment, it will retain many of the central classical principles, including non-contradiction, excluded middle, the introduction and elimination rules for conjunction and disjunction, and De Morgan's laws, all of which have been invalidated by various theories of epistemic modals. Moreover, our logic will, uniquely among E-logics, be fully classical over parts of the modal fragment restricted to the same "epistemic level"—essentially, fragments built up out of sentences with uniform levels of nesting of modal operators—limiting non-classicality to the part of the modal fragment that involves combinations of sentences across different epistemic levels.

To be clear, we do not have an argument that the logic we present is *the* minimal variation on classical logic that is an E-logic and invalidates the patterns that intuitively fail for epistemic modals. There are E-logics which validate strictly more classical patterns than ours; and there may turn out to be a case for adopting some such extension of our logic (as we will briefly discuss in § 5.1).[6] But as we will discuss in § 6, our approach preserves a lot more of classical logic than any other E-logic we know of in the literature. Among other things, we hope that our discussion will spur exploration of other E-logics extending ours.

3 Algebraic semantics

In this section, we begin our development of formal semantics for epistemic modals. We start with a rather abstract *algebraic semantics*. A model in such a semantics associates with each formula of the formal language an element in an *algebra of propositions*. This is familiar from *possible world semantics* for classical propositional logic, wherein a model associates with each formula a set of possible worlds, i.e., an element of the powerset algebra $\mathcal{P}(W)$ arising from the set W of worlds. The meanings of 'and', 'or', and 'not' are given by the operations of intersection, union, and complementation, respectively, in the powerset algebra $\mathcal{P}(W)$. The only difference between possible world semantics and algebraic semantics for classical propositional logic is that in the latter, we allow Boolean algebras other than powerset algebras. Instead of associating each formula with an element of $\mathcal{P}(W)$ for some set W, we associate with each formula an element of an arbitrary Boolean algebra B, which comes equipped with operations $\neg$, $\wedge$, and $\vee$ used to interpret 'not', 'and', and 'or' (see § 3.1). The powerset algebra $\mathcal{P}(W)$ is just one concrete example of a Boolean algebra. But for the purposes of classical propositional logic, there is no loss of generality in working only with powerset algebras.

Possible world semantics for normal modal logic can also be viewed as a concrete version of an algebraic semantics. A set W together with a binary accessibility relation R on W gives us not only the powerset

[6] In terms defined in Section 3, there are orthologics stronger than the minimal orthologic in which orthomodularity is still not derivable (see Harding 1988), and such logics can be extended to epistemic orthologics in the sense of Definition 3.24. Though we have arguments against some of these stronger logics, we have no sweeping argument against all stronger logics. In § 5.1, we raise the possibility of strengthening our logic with modal principles.

algebra $\mathcal{P}(W)$ but also a modal operator $\Diamond$ on $\mathcal{P}(W)$ defined for $A \in \mathcal{P}(W)$ by

$$\Diamond A = \{w \in W \mid \exists v : wRv \text{ and } v \in A\}.$$

A more abstract algebraic semantics uses an arbitrary Boolean algebra equipped with a modal operator, called a *Boolean algebra with operator* (BAO). In the context of modal logic, there *is* a loss of generality in working only with powerset algebras, as not all normal modal logics can be given possible world semantics as above (see Holliday and Litak 2019 and references therein). But the normal modal logics discussed in connection with natural language semantics typically can be handled by possible world semantics.

The logic we are after in this paper is *non-classical*. Thus, we cannot use powerset algebras, as in possible world semantics, or Boolean algebras more generally, as in algebraic semantics for classical logic. Instead, we will use a more general class of algebras, known as *ortholattices*, in which it is possible to invalidate classical laws such as distributivity. We will add modal operators to ortholattices to give an algebraic semantics for an *epistemic orthologic*. Then in § 4, we will give a more concrete semantics, known as *possibility semantics*, that stands to ortholattice semantics as possible world semantics stands to Boolean algebraic semantics.

3.1 Review of algebraic semantics for orthologic

In this section, we review ortholattices and their associated *orthologic*. Ortholattices have long been important objects of study in lattice theory (see Birkhoff 1967, § 14), examples of which arise as algebras of events in quantum mechanics (Birkhoff and von Neumann 1936). We will argue that ortholattices also arise as algebras of propositions in a language with epistemic modals. Those familiar with basic lattice theory and ortholattices may skip straight to § 3.2, where we add epistemic modals to the picture.

3.1.1 Ortholattices

First, we recall the definition of a lattice, where we think of A as a set of propositions and $\vee$ and $\wedge$ as the operations of disjunction and conjunction, respectively.

Definition 3.1. A *lattice* is a tuple $L = \langle A, \vee, \wedge \rangle$ where A is a nonempty set and $\vee$ and $\wedge$ are binary operations on A such that the following equations hold for all $a, b, c \in A$ and $\circ \in \{\vee, \wedge\}$:

- idempotence: $a \circ a = a$;
- associativity: $a \circ (b \circ c) = (a \circ b) \circ c$;
- commutativity: $a \circ b = b \circ a$;
- absorption: $a \wedge (a \vee b) = a \vee (a \wedge b) = a$.

We define a binary relation $\leq$ on A, called the *lattice order of* L, by: $a \leq b$ if and only if $a = a \wedge b$.

It is easy to check that $a \leq b$ is also equivalent to $a \vee b = b$ (using absorption and commutativity). Moreover, $\leq$ is a partial order (i.e., reflexive, transitive, and antisymmetric). Indeed, let us recall the order-theoretic definition of a lattice as a partially ordered set. Given a partial order $\leqslant$ on a set A,

- an *upper bound* of a subset $X \subseteq A$ is a $y \in A$ such that $x \leqslant y$ for every $x \in X$;
- a *least upper bound* of X is an upper bound y of X such that $y \leqslant z$ for every upper bound z of X.

If there is a least upper bound of X, it is unique by the antisymmetry of $\leqslant$. The notion of a *greatest lower bound* is defined dually. Then a partially ordered set is a lattice if every nonempty finite subset has both a least upper bound and a greatest lower bound. From such a partially ordered set, we obtain a lattice

$\langle A, \vee, \wedge \rangle$ in the sense of Definition 3.1 by defining $a \vee b$ to be the least upper bound of $\{a, b\}$ and $a \wedge b$ to be the greatest lower bound of $\{a, b\}$. Conversely, given a lattice in the sense of Definition 3.1, the partially ordered set $\langle A, \leq \rangle$ with $\leq$ defined as in Definition 3.1 is a lattice in the order-theoretic sense. Hence we may think of lattices in terms of either the equational or order-theoretic definition. Finally, a lattice is *complete* if every subset has both a least upper bound and a greatest lower bound. Every finite lattice is complete but there are infinite lattices that are not complete.

We will assume that our lattices have bounds, corresponding to the contradictory proposition $\bot$ and the trivial proposition $\top$.

Definition 3.2. A *bounded lattice* is a tuple $L = \langle A, \vee, 0, \wedge, 1 \rangle$ where $\langle A, \vee, \wedge \rangle$ is a lattice and 0 and 1 are elements of A such that for all $a \in A$, we have

- boundedness: $a \vee 0 = a$ and $a \wedge 1 = a$.

Order-theoretically, a lattice is bounded if its order has a minimum element and a maximum element.

Adding an operation $\neg$ for negation finally brings us to the definition of an ortholatticce.

Definition 3.3. An *ortholattice* is a tuple $\langle A, \vee, 0, \wedge, 1, \neg \rangle$ where $\langle A, \vee, 0, \wedge, 1 \rangle$ is a bounded lattice and $\neg$ is a unary operation on A, called the *orthocomplementation*, that satisfies:

1. complementation: for all $a \in A$, $a \vee \neg a = 1$ and $a \wedge \neg a = 0$;
2. involution: for all $a \in A$, $\neg\neg a = a$;
3. order-reversal: for all $a, b \in A$, if $a \leq b$, then $\neg b \leq \neg a$.

An equivalent definition replaces 3 with (either one of) *De Morgan's laws*:

- for all $a, b \in A$, $\neg(a \vee b) = \neg a \wedge \neg b$;
- for all $a, b \in A$, $\neg(a \wedge b) = \neg a \vee \neg b$.

The difference between arbitrary ortholattices and the Boolean algebras of classical logic is that ortholattices need not obey the distributive law, as we will see in Examples 3.8 and 3.20.

Definition 3.4. A *Boolean algebra* is a tuple $\langle A, \vee, 0, \wedge, 1, \neg \rangle$ where $\langle A, \vee, 0, \wedge, 1 \rangle$ is a bounded lattice, $\neg$ is a unary operation on A satisfying complementation as in Definition 3.3.1, and the *distributive laws* hold:

- for all $a, b, c \in A$, $a \wedge (b \vee c) = (a \wedge b) \vee (a \wedge c)$;
- for all $a, b, c \in A$, $a \vee (b \wedge c) = (a \vee b) \wedge (a \vee c)$.[7]

It is straightforward to prove that every Boolean algebra is an ortholattice.

The following weakening of distributivity is important in the study of ortholattices arising in quantum mechanics (see, e.g., Chiara and Giuntini 2002).

Definition 3.5. An *orthomodular lattice* is an ortholattice satisfying the orthomodular law:

- for all $a, b \in A$, $a \vee (\neg a \wedge (a \vee b)) = a \vee b$.

Or equivalently, for all $a, c \in A$, if $a \leq c$, then $c \leq a \vee (\neg a \wedge c)$ (equivalently, $c = a \vee (\neg a \wedge c)$).

[7]In fact, a lattice satisfies the first bullet point if and only if it satisfies the second, so including both is redundant.

Associated with the difference between Boolean algebras and arbitrary ortholattices concerning distributivity is a difference concerning the interaction of conjunction, contradiction, and negation, given by the following standard fact.

Lemma 3.6. In a Boolean algebra, $\neg$ is *pseudocomplementation*: for all $a, b \in A$, $a \wedge b = 0$ implies $b \leq \neg a$, so $\neg a$ is the greatest element with respect to $\leq$ of the set $\{b \in A \mid a \wedge b = 0\}$.[8]

Crucially, the orthocomplementation in an ortholattice is not necessarily pseudocomplementation. In fact, the orthocomplementation being pseudocomplementation implies that the ortholattice is Boolean.

Proposition 3.7. For any ortholattice L, the following are equivalent:

1. L is a Boolean algebra;
2. L is distributive;
3. the orthocomplementation operation in L is pseudocomplementation.

Proof. The equivalence of 1 and 2 is straightforward, as noted above.

From 2 to 3, we show that distributivity implies pseudocomplementation over ortholattices. Suppose $a \wedge b = 0$, so $a \wedge b \leq \neg b$. Then since $a \wedge \neg b \leq \neg b$, we have $(a \wedge b) \vee (a \vee \neg b) \leq \neg b$. Finally, by distributivity, $a \wedge (b \vee \neg b) \leq (a \wedge b) \vee (a \vee \neg b)$, so $a \leq \neg b$.

From 3 to 2, we show that pseudocomplementation implies distributivity over ortholattices. First note that pseudocomplementation implies disjunctive syllogism: $(x \vee y) \wedge \neg x \leq y$. For $(x \vee y) \wedge \neg x \wedge \neg y \leq 0$ by De Morgan's laws and complementation, so $(x \vee y) \wedge \neg x \leq y$ by pseudocomplementation and involution. Now for distributivity, we have

$$
\begin{aligned}
& a \wedge (b \vee c) \wedge (\neg a \vee \neg b) \leq a \wedge (b \vee c) \wedge \neg b \leq a \wedge c \quad \text{using disjunctive syllogism twice} \\
\Rightarrow\ & a \wedge (b \vee c) \wedge (\neg a \vee \neg b) \wedge \neg(a \wedge c) \leq \neg(a \wedge c) \wedge (a \wedge c) \leq 0 \\
\Rightarrow\ & a \wedge (b \vee c) \leq \neg((\neg a \vee \neg b) \wedge \neg(a \wedge c)) \quad \text{by pseudocomplementation} \\
\Rightarrow\ & a \wedge (b \vee c) \leq (a \wedge b) \vee (a \wedge c) \quad \text{by De Morgan's and involution.} \quad \square
\end{aligned}
$$

Thus, we obtain a three-way equivalence for any ortholattice L between being a Boolean algebra, being distributive, and having its orthocomplementation be pseudocomplementation.

Example 3.8. Figure 1 shows Hasse diagrams of the ortholattices $\mathbf{O}_6$ and $\mathbf{MO}_2$. Recall that in a Hasse diagram of a lattice with lattice order $\leq$, a line segment going *upward* from x to y means that $x \leq y$ and there is no third element z with $x \leq z \leq y$. Observe that $\mathbf{O}_6$ is not orthomodular and hence not distributive. For $a \leq b$ and yet $a \vee (\neg a \wedge b) = a \vee 0 = a \neq b$. Also note that the orthocomplementation is not pseudocomplementation: $\neg a \wedge b = 0$ and yet $b \nleq \neg\neg a = a$. $\mathbf{MO}_2$ is orthomodular[9] but it is not distributive: $a \wedge (\neg a \vee b) = a \wedge 1 = a \neq 0 = 0 \vee 0 = (a \wedge \neg a) \vee (a \wedge b)$. Also note that orthocomplementation is not pseudocomplementation: $a \wedge b = 0$ and yet $b \nleq \neg a$.

3.1.2 Language and consequence

We can use ortholattices to interpret a basic propositional logical language.

[8] This fact also holds for Heyting algebras, which provide algebraic semantics for *intuitionistic logic*, where $\neg a := a \to 0$.

[9] In fact, it is a *modular* lattice, meaning, again, that for all $a, b, c \in A$, if $a \leq c$, then $a \vee (b \wedge c) = (a \vee b) \wedge c$.

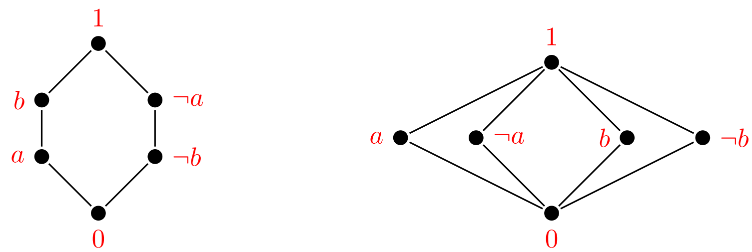

Figure 1: Hasse diagrams of the ortholattices $\mathbf{O}_6$ (left) and $\mathbf{MO}_2$ (right).

Definition 3.9. Let $\mathcal{L}$ be the language generated by the grammar

$$\varphi ::= \top \mid p \mid \neg\varphi \mid (\varphi \wedge \varphi)$$

where p belongs to a countably infinite set Prop of propositional variables.

We define $\varphi \vee \psi := \neg(\neg\varphi \wedge \neg\psi)$, a definition justified by the fact that ortholattices satisfy De Morgan's laws and involution. Note that we use the same symbols for the connectives of $\mathcal{L}$ and the operations in ortholattices, trusting that no confusion will arise.

As usual in algebraic semantics, we interpret propositional variables as elements of an algebra and then extend the interpretation to all formulas of the language recursively.

Definition 3.10. A *valuation* on an ortholattice $\langle A, \vee, 0, \wedge, 1, \neg \rangle$ is a map $\theta : \mathsf{Prop} \to A$. Such a θ extends to $\tilde{\theta} : \mathcal{L} \to A$ by: $\tilde{\theta}(\top) = 1$, $\tilde{\theta}(\neg\varphi) = \neg\tilde{\theta}(\varphi)$, and $\tilde{\theta}(\varphi \wedge \psi) = \tilde{\theta}(\varphi) \wedge \tilde{\theta}(\psi)$.

Also as usual, we say that ψ is a semantic consequence of φ if the semantic value of φ is always below the semantic value of ψ in the lattice order $\leq$ (which one may now view as an entailment relation, just as the subset relation is viewed as an entailment relation between propositions in possible world semantics).

Definition 3.11. Given a class $\mathbf{C}$ of ortholattices, we define the semantic consequence relation $\vDash_{\mathbf{C}}$, a binary relation on $\mathcal{L}$, as follows: $\varphi \vDash_{\mathbf{C}} \psi$ if for all $L \in \mathbf{C}$ and valuations θ on L, we have $\tilde{\theta}(\varphi) \leq \tilde{\theta}(\psi)$, where $\leq$ is the lattice order of L.

We can similarly define a consequence relation between a set of premises on the left and a single conclusion on the right: $\Gamma \vDash_{\mathbf{C}} \psi$ if for all $L \in \mathbf{C}$, valuations θ on L, and $a \in L$, if a is a lower bound of $\{\tilde{\theta}(\varphi) \mid \varphi \in \Gamma\}$, then $a \leq \tilde{\theta}(\psi)$. But for simplicity we will only consider finite sets of premises here, in which case a single premise on the left suffices given that $\mathcal{L}$ contains a conjunction interpreted as meet.

3.1.3 Logic, soundness and completeness

Axiomatizing the semantic consequence relation of Definition 3.11 is straightforward.

Definition 3.12 (Goldblatt 1974)**.** An *orthologic* is a binary relation $\vdash$ on the set $\mathcal{L}$ of formulas such that for all $\varphi, \psi, \chi \in \mathcal{L}$:

1. $\varphi \vdash \top$;
2. $\varphi \vdash \varphi$;
3. $\varphi \wedge \psi \vdash \varphi$;
4. $\varphi \wedge \psi \vdash \psi$;
5. $\varphi \vdash \neg\neg\varphi$;
6. $\neg\neg\varphi \vdash \varphi$;
7. $\varphi \wedge \neg\varphi \vdash \psi$;
8. if $\varphi \vdash \psi$ and $\psi \vdash \chi$, then $\varphi \vdash \chi$;
9. if $\varphi \vdash \psi$ and $\varphi \vdash \chi$, then $\varphi \vdash \psi \wedge \chi$;
10. if $\varphi \vdash \psi$, then $\neg\psi \vdash \neg\varphi$.

A *theorem* of the orthologic $\vdash$ is a formula φ such that $\top \vdash \varphi$. As the intersection of orthologics is clearly an orthologic, there is a smallest orthologic, denoted O or $\vdash_{\mathsf{O}}$.

Note that with $\varphi \vee \psi$ defined as $\neg(\neg\varphi \wedge \neg\psi)$, we get the introduction and elimination rules for disjunction:

- $\varphi \vdash \varphi \vee \psi$.
- if $\varphi \vdash \chi$ and $\psi \vdash \chi$, then $\varphi \vee \psi \vdash \chi$.

For the first, by rule 3, we have $\neg\varphi \wedge \neg\psi \vdash \neg\varphi$, so by rule 10, we have $\neg\neg\varphi \vdash \varphi \vee \psi$, which with rules 5 and 8 yields $\varphi \vdash \varphi \vee \psi$. For the second, if $\varphi \vdash \chi$ and $\psi \vdash \chi$, then $\neg\chi \vdash \neg\varphi$ and $\neg\chi \vdash \neg\psi$ by rule 10, which implies $\neg\chi \vdash \neg\varphi \wedge \neg\psi$ by rule 9, which implies $\varphi \vee \psi \vdash \neg\neg\chi$ by rule 10 and hence $\varphi \vee \psi \vdash \chi$ by rules 6 and 8.

However, we do *not* have pseudocomplementation or distributivity, either of which would collapse O to classical logic (recall Proposition 3.7). Nor do we have what might be called *proof by cases with side assumptions*:

- if $\xi \wedge \varphi \vdash \chi$ and $\xi \wedge \psi \vdash \chi$, then $\xi \wedge (\varphi \vee \psi) \vdash \chi$.

This would allow the derivation of distributivity, since disjunction introduction and proof by cases with side assumptions yield

$$\xi \wedge \varphi \vdash (\xi \wedge \varphi) \vee (\xi \wedge \psi) \text{ and } \xi \wedge \psi \vdash (\xi \wedge \varphi) \vee (\xi \wedge \psi), \text{ so } \xi \wedge (\varphi \vee \psi) \vdash (\xi \wedge \varphi) \vee (\xi \wedge \psi).$$

A completeness theorem can now be proved using standard techniques of algebraic logic.

Theorem 3.13. The logic O is sound and complete with respect to the class $\mathbf{O}$ of all ortholattices according to the consequence relation of Definition 3.11: for all $\varphi, \psi \in \mathcal{L}$, we have $\varphi \vdash_{\mathsf{O}} \psi$ if and only if $\varphi \vDash_{\mathbf{O}} \psi$.

Proof. For soundness, it is easy to check that $\varphi \vdash_{\mathsf{O}} \psi$ implies $\varphi \vDash_{\mathbf{O}} \psi$. For completeness, recall the construction of the *Lindenbaum-Tarski algebra L of* O: the underlying set of L is the set of all equivalence classes of formulas of $\mathcal{L}$, where φ and ψ are equivalent if $\varphi \vdash \psi$ and $\psi \vdash \varphi$; then let $0 = [\neg\top]$ and $1 = [\top]$, and given equivalence classes $[\varphi]$ and $[\psi]$, define $[\varphi] \vee [\psi] = [\varphi \vee \psi]$, $[\varphi] \wedge [\psi] = [\varphi \wedge \psi]$, $\neg[\varphi] = [\neg\varphi]$. Then where $\leq$ is the lattice order of L, we have $[\varphi] \leq [\psi]$ iff $\varphi \vdash_{\mathsf{O}} \psi$. It is easy to check that L is an ortholattice. Moreover, where θ is the valuation with $\theta(p) = [p]$ for each $p \in \mathsf{Prop}$, an obvious induction shows $\tilde{\theta}(\varphi) = [\varphi]$ for each $\varphi \in \mathcal{L}$. Now if $\varphi \nvdash_{\mathsf{O}} \psi$, then $[\varphi] \not\leq [\psi]$ and hence $\tilde{\theta}(\varphi) \not\leq \tilde{\theta}(\psi)$, so $\varphi \nvDash_{\mathbf{O}} \psi$. $\square$

Remark 3.14. A Fitch-style natural deduction system for the minimal orthologic in the signature $\top, \neg, \wedge, \vee$ can be obtained from a Fitch-style natural deduction system for classical logic in the same signature (cf. Fitch 1952, 1966) by (i) restricting the rule of *Reiteration* so that the formula occurrence to be reiterated and its reiterates must belong to the same column and the same subproofs, (ii) keeping the introduction and elimination rules for $\wedge$ and $\vee$ the same,[10] and (iii) allowing $\neg$ Introduction and Reductio ad Absurdum to apply when there is a contradiction between a formula derived in a subproof and a formula previously derived in the column in which the subproof immediately occurs. Such a proof system is shown in Figure 2. The motivation for restricting Reiteration can easily be seen when we consider epistemic modal propositions, to

[10]We have in mind formulations of the introduction and elimination rules for $\wedge$ and $\vee$, as in Fitch 1952, 1966 (but unlike some introductory logic texts), that do not themselves build in the power to draw previous formulas into subproofs; the Reiteration rule must always be used for this purpose.

be added in the next section, as in the following example:

1	$\Diamond p$	
2	$(p \vee \neg p)$	
3	p	
4	$p \vee (\neg p \wedge \Diamond p)$	$\vee$I, 3
5	$\neg p$	
6	$\Diamond p$	Reit, 1
7	$\neg p \wedge \Diamond p$	$\wedge$I, 5, 6
8	$p \vee (\neg p \wedge \Diamond p)$	$\vee$I, 7
9	$p \vee (\neg p \wedge \Diamond p)$	$\vee$E, 2, 3–4, 5–8

The application of Reiteration on line 6 is obviously suspect: we should not be allowed to reiterate the assumption that *might* p into a subproof where we have just supposed *not* p! Moreover, if q is a genuine propositional variable, standing in for any proposition (cf. § 3.2.4), including an epistemic proposition such as *might* p, then we cannot accept the analogous proof with q in place of $\Diamond p$. However, if we restrict the Reiteration rule as suggested above, then all is well: where $\varphi \vdash_{\mathsf{FitchO}} \psi$ means that there is a Fitch-style proof of the conclusion ψ from the assumption φ, Theorem 3.13 holds with $\vdash_{\mathsf{FitchO}}$ in place of $\vdash_{\mathsf{O}}$. For soundness, an easy induction shows that if there is a subproof or proof beginning with φ and ending with ψ, then—thanks to the restriction on Reiteration—ψ is in fact a semantic consequence of φ, i.e., $\varphi \vDash_{\mathbf{O}} \psi$. For completeness, it is easy to check that the principles of Definition 3.12, as well as the equivalence of $\varphi \vee \psi$ and $\neg(\neg\varphi \wedge \neg\psi)$, hold for $\vdash_{\mathsf{FitchO}}$, in which case the same style of proof as for Theorem 3.13 applies.

3.2 Adding epistemic modality

In this section, we extend the algebraic semantics of § 3.1 to interpret the modals 'must' and 'might'.

3.2.1 Modal and epistemic ortholattices

We begin by extending ortholattices with a unary operation $\Box$. First, we define a baseline notion of a *modal* ortholattice and then add additional conditions for epistemic ortholattices.

Definition 3.15. A *modal ortholattice* is a tuple $\langle A, \vee, 0, \wedge, 1, \neg, \Box\rangle$ where $\langle A, \vee, 0, \wedge, 1, \neg\rangle$ is an ortholattice and $\Box$ is a unary operation on A satisfying:

- $\Box(a \wedge b) = \Box a \wedge \Box b$ for all $a, b \in A$, and $\Box 1 = 1$.

For $a \in A$, we define $\Diamond a = \neg\Box\neg a$.

The following is easy to check.

Lemma 3.16. In any modal ortholattice, we have:

- $\Diamond(a \vee b) = \Diamond a \vee \Diamond b$ for all $a, b \in A$, and $\Diamond 0 = 0$.

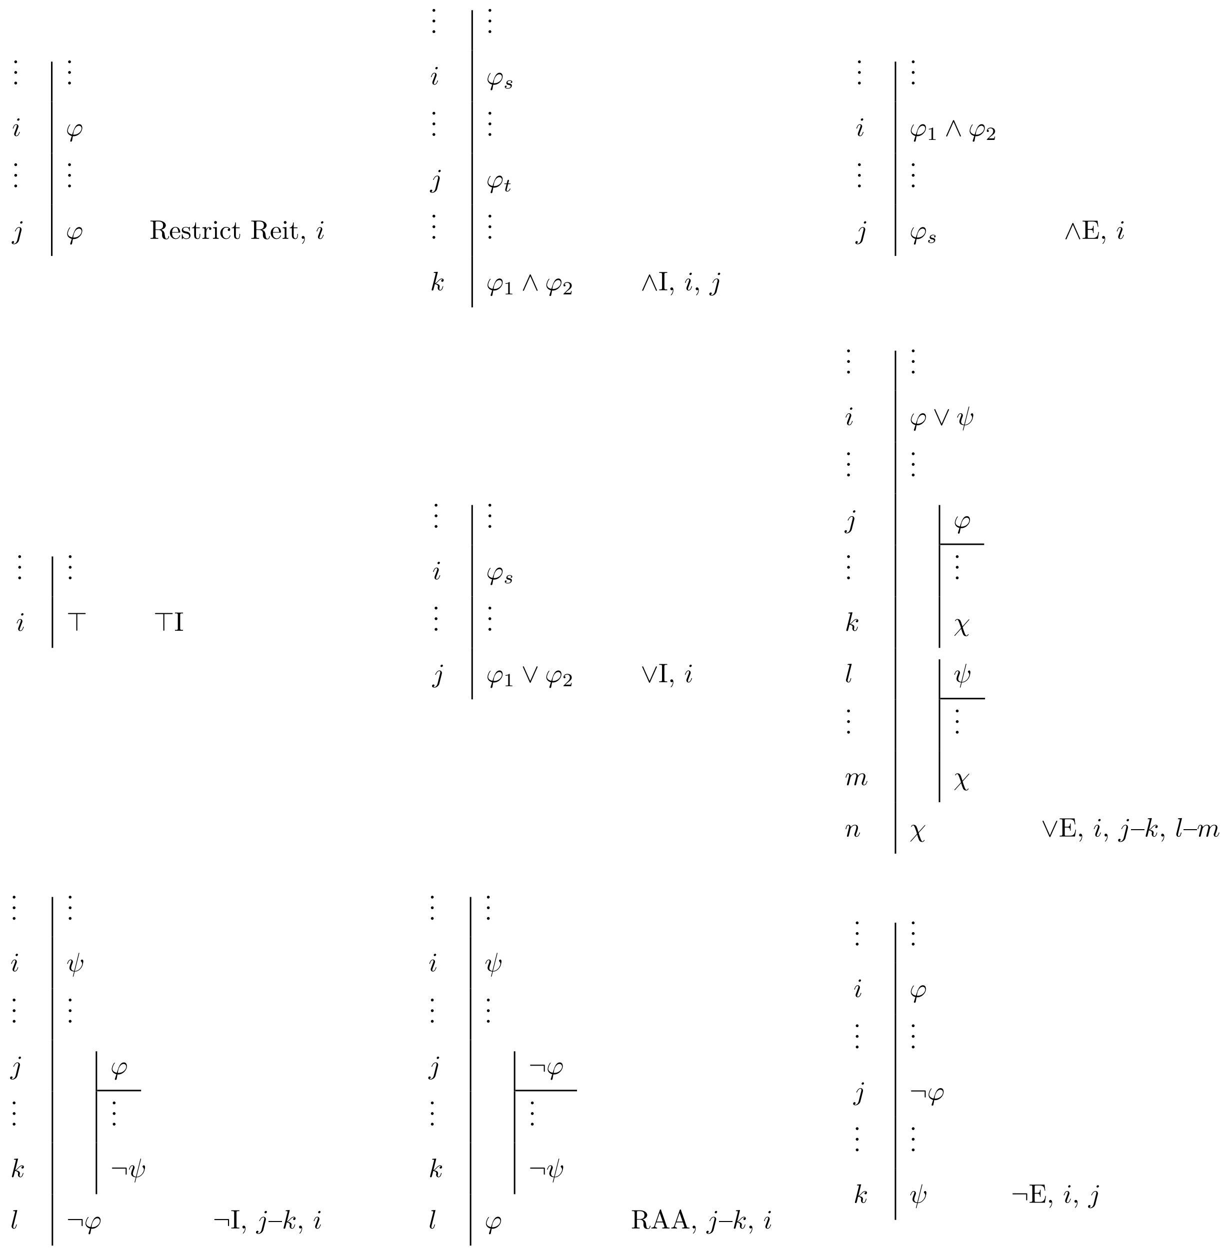

Figure 2: Rules of a Fitch-style proof system for the minimal orthologic, where $s, t \in \{1, 2\}$.

Indeed, we could have defined modal ortholattices as algebras $\langle A, \vee, 0, \wedge, 1, \neg, \Diamond \rangle$ where $\langle A, \vee, 0, \wedge, 1, \neg \rangle$ is an ortholattice and $\Diamond$ is a unary operation on A satisfying the conditions of Lemma 3.16. But later it will turn out to be more convenient to have $\Box$ as our primitive.

Now we can consider additional constraints on the $\Box$ operation, just as modal logic considers additional axioms on a $\Box$ modality (see, e.g., Chagrov and Zakharyaschev 1997, p. 116). In particular, in order to view $\Box$ and $\Diamond$ as 'must' and 'might', we adopt two further constraints. First, corresponding to the factivity of 'must', we have the following constraint.

Definition 3.17. A T *modal ortholattice* is a modal ortholattice also satisfying

- $\Box a \leq a$ for all $a \in A$.

Next comes the crucial constraint corresponding to Wittgenstein sentences being contradictions.

Definition 3.18. An *epistemic ortholattice* is a T modal ortholattice also satisfying

- Wittgenstein's Law: $\neg a \wedge \Diamond a = 0$ for all $a \in A$.

By the involution property of $\neg$ in an ortholattice, Wittgenstein's Law is equivalent to $a \wedge \Diamond \neg a = 0$; and by the commutativity of $\wedge$ in a lattice, it is also equivalent to $\Diamond a \wedge \neg a = 0$ (and $\Diamond \neg a \wedge a = 0$), in contrast to the consistency of $\Diamond a \wedge \neg a$ (and $\Diamond \neg a \wedge a$) in some dynamic systems (e.g. in Groenendijk et al. 1996; see van Benthem 1996).

Some philosophers of language have argued for *iteration principles* for 'must' and 'might', leading to the following additional constraints.

Definition 3.19. An $\mathsf{S5}$ *modal ortholattice* is a T modal ortholattice also satisfying:

- $\Box a \leq \Box\Box a$ for all $a \in A$;
- $\Diamond a \leq \Box \Diamond a$ for all $a \in A$.

A $\mathsf{S5}$ *epistemic ortholattice* is an $\mathsf{S5}$ modal ortholattice that is also an epistemic ortholattice.

Example 3.20. Figure 3 displays the Hasse diagram of an $\mathsf{S5}$ epistemic ortholattice L, labeled in two ways. Recall that a line segment going *upward* from x to y means that $x \leq y$ and there is no third element z with $x \leq z \leq y$. The $\Box$ operation is depicted by the blue arrows. Note the failure of distributivity:

$$(p \vee \neg p) \wedge (\Diamond p \wedge \Diamond \neg p) = 1 \wedge (\Diamond p \wedge \Diamond \neg p) = \Diamond p \wedge \Diamond \neg p \neq 0$$

and yet

$$(p \wedge \Diamond \neg p) \vee (\neg p \wedge \Diamond p) = 0 \vee 0 = 0.$$

That distributivity fails follows from the failure of orthomodularity. Recall this is the condition that $a \leq b$ implies $a \vee (\neg a \wedge b) = b$. Yet in Figure 3 we have $p \leq \Diamond p$ and yet $p \vee (\neg p \wedge \Diamond p) = p \vee 0 = p \neq \Diamond p$.

Next, observe that

$$p \wedge \Diamond \neg p = 0 \text{ and yet } \Diamond \neg p \nleq \neg p.$$

This shows that the orthocomplementation $\neg$ is not *pseudocomplementation* (recall Lemma 3.6).

Also observe that

$$(p \vee \Box \neg p) \wedge \Diamond p = \top \wedge \Diamond p = \Diamond p \nleq p.$$

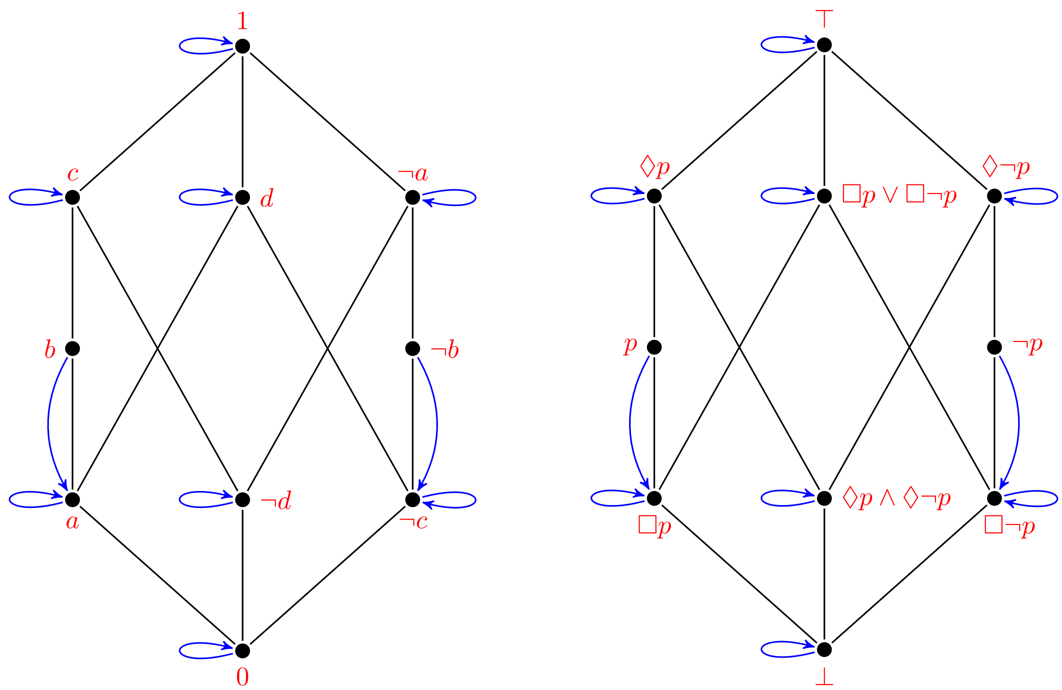

Figure 3: Hasse diagram of an S5 epistemic ortholattice with two labelings

This shows that disjunctive syllogism fails.

Finally, note that the *subortholattice* of L generated by p is the four-element *Boolean* algebra:

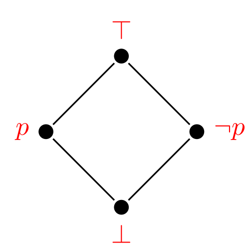

Thus, if we think of p as a non-modal proposition such as 'it is raining' and generate further propositions using disjunction, conjunction, and negation, the result is Boolean. This corresponds to the fact that the failures of distributivity, orthomodularity, pseudocomplementation, and disjunctive syllogism discussed above essentially involve epistemic modals. In § 3.2.4, we return to this observation and show how to recover full classical reasoning for a Boolean fragment of our language interpreted in Boolean subalgebras of our lattices.

3.2.2 Language and consequence

We can use epistemic ortholattices to interpret the following propositional modal language.

Definition 3.21. Let $\mathcal{EL}$ be the language generated by the grammar

$$\varphi ::= \top \mid p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \Box\varphi$$

where p belongs to a countably infinite set Prop of propositional variables.

We define $\bot := \neg\top$, $\varphi \vee \psi := \neg(\neg\varphi \wedge \neg\psi)$, and $\Diamond\varphi := \neg\Box\neg\varphi$. Again we use the same symbols for the connectives of $\mathcal{EL}$ and the operations in epistemic ortholattices, trusting that no confusion will arise.

The algebraic semantics follows the same approach as in § 3.1.2.

Definition 3.22. A *valuation* on a modal ortholattice $\langle A, \vee, 0, \wedge, 1, \neg, \Box\rangle$ is a map $\theta : \mathsf{Prop} \to A$. Such a θ extends to $\tilde{\theta} : \mathcal{EL} \to A$ by: $\tilde{\theta}(\top) = 1$, $\tilde{\theta}(\neg\varphi) = \neg\tilde{\theta}(\varphi)$, $\tilde{\theta}(\varphi \wedge \psi) = \tilde{\theta}(\varphi) \wedge \tilde{\theta}(\psi)$, and $\tilde{\theta}(\Box\varphi) = \Box\tilde{\theta}(\varphi)$.

Definition 3.23. Given a class $\mathbf{C}$ of modal ortholattices, define the semantic consequence relation $\vDash_{\mathbf{C}}$, a binary relation on $\mathcal{EL}$, as follows: $\varphi \vDash_{\mathbf{C}} \psi$ if for every $L \in \mathbf{C}$ and valuation θ on L, we have $\tilde{\theta}(\varphi) \leq \tilde{\theta}(\psi)$, where $\leq$ is the lattice order of L.

3.2.3 Logic, soundness and completeness

Building on § 3.1.3, axiomatizing the semantic consequence relation of Definition 3.23 is straightforward.

Definition 3.24. An *epistemic orthologic* is a binary relation $\vdash$ on the set $\mathcal{EL}$ of formulas satisfying for all $\varphi, \psi \in \mathcal{EL}$ conditions 1-10 of Definition 3.12 plus:

11. if $\varphi \vdash \psi$, then $\Box\varphi \vdash \Box\psi$;
12. $\Box\varphi \wedge \Box\psi \vdash \Box(\varphi \wedge \psi)$;
13. $\varphi \vdash \Box\top$;
14. $\Box\varphi \vdash \varphi$;
15. $\neg\varphi \wedge \Diamond\varphi \vdash \bot$ (Wittgenstein's Law).

As the intersection of epistemic orthologics is clearly an epistemic orthologic, there is a smallest epistemic orthologic, denoted EO or $\vdash_{\mathsf{EO}}$.[11]

Wittgenstein's Law yields the following noteworthy property of epistemic orthologics.

Lemma 3.25. For any epistemic orthologic $\vdash$ and $\varphi \in \mathcal{EL}$, if $\Box\varphi \vdash \bot$, then $\varphi \vdash \bot$.

Proof. If $\Box\varphi \vdash \bot$, then $\top \vdash \neg\Box\varphi$, in which case $\varphi \vdash \varphi \wedge \neg\Box\varphi \vdash \neg\neg\varphi \wedge \Diamond\neg\varphi \vdash \bot$, so $\varphi \vdash \bot$. ☐

Conversely, if one assumes the principle in Lemma 3.25 together with principles 11 and 14 in Definition 3.24 on top of orthologic, then Wittgenstein's Law is derivable (see the proof of Fact 3.28).

As in § 3.1.3, a completeness theorem can be proved using standard techniques of algebraic logic.

Theorem 3.26. The logic EO is sound and complete with respect to the class $\mathbf{EO}$ of epistemic ortholattices according to the consequence relation of Definition 3.23: for all $\varphi, \psi \in \mathcal{EL}$, we have $\varphi \vdash_{\mathsf{EO}} \psi$ if and only if $\varphi \vDash_{\mathbf{EO}} \psi$.

Proof. By the same strategy as in the proof of Theorem 3.13, checking that the Lindenbaum-Tarski algebra of EO is an epistemic ortholattice. ☐

Remark 3.27. Soundness and completeness with respect to $\mathsf{S5}$ epistemic ortholattices is also straightforward by adding the rules $\Box\varphi \vdash \Box\Box\varphi$ (known as $\mathsf{4}$) and $\Diamond\varphi \vdash \Box\Diamond\varphi$ (known as $\mathsf{5}$) to Definition 3.24. However, neither of these rules, nor the rule $\varphi \vdash \Box\Diamond\varphi$ (known as B), is valid with respect to epistemic ortholattices in general. Moreover, the evidence about the status of these inferences for epistemic modality is mixed.

[11] A Fitch-style proof system for EO can be obtained from the Fitch-style proof system for O in Figure 2 by adding rules of $\Box$ introduction, $\Box$ elimination, and Strict Reiteration as in Fitch 1966, plus an additional rule of Epistemic Contradiction that is like the $\neg$E rule in Figure 2 but with $\neg\varphi$ replaced by $\neg\Box\varphi$.

On the one hand, Moss (2015) argues that patterns from nested epistemic modals tell against the collapse principles that these axioms together entail. For instance, the sentences in (22) do not sound obviously equivalent but instead are intuitively increasingly strong:[12]

(22) a. John might possibly win.
b. John might win.
c. John certainly might win.

Judgments here are somewhat difficult to ascertain, however, given that it is very difficult to directly stack epistemic modals in English.

Pulling in the other direction, conjunctions that instantiate violations of 4, 5, and B—that is, sentences with the form $\Box p \wedge \neg\Box\Box p$, $\Diamond p \wedge \neg\Box\Diamond p$, and $p \wedge \neg\Box\Diamond p$, respectively—do sound inconsistent, as in (23) (using duality to improve readability):

(23) a. #John must be the winner, but maybe he might not be the winner.
b. #John might be the winner, but it might be that he must not be the winner.
c. #John is the winner, but it might be that he must not be the winner.

This extends to embedded environments, suggesting this is not simply a Moorean phenomenon but should be accounted for logically.

In fact, the account so far has a nice way to make sense of both of these kinds of evidence. On the one hand, 4, 5, and B are all invalid according to our consequence relation for **EO**. On the other hand, the conjunctions above that instantiate their violations are all inconsistent. Reasoning from that inconsistency to the validity of 4, 5, and B, although classically valid, is blocked in our system since negation is not pseudocomplementation. This lets us account both for the fact that nested modals appear not to collapse and for the fact that conjunctions that witness the failures of 4, 5, and B appear inconsistent.

One might think that without 4, which is equivalent to $\Diamond\Diamond\varphi \vdash \Diamond\varphi$, the other principles of our logic could not account for the fact that $p \wedge \Diamond\Diamond\neg p$ should be inconsistent. In fact, we do not need 4 to account for this. For the following, let $\Diamond^0\varphi = \varphi$ and $\Diamond^{n+1}\varphi = \Diamond\Diamond^n\varphi$.

Fact 3.28. For any epistemic orthologic $\vdash$, $n \in \mathbb{N}$, and $\varphi \in \mathcal{EL}$, we have $\varphi \wedge \Diamond^n\neg\varphi \vdash \bot$.

Proof. By induction on n. Assume $\psi \wedge \Diamond^n\neg\psi \vdash \bot$ for all $\psi \in \mathcal{EL}$. Then for any $\varphi \in \mathcal{EL}$,

$$\Box(\varphi \wedge \Diamond^{n+1}\neg\varphi) \vdash \Box\varphi \wedge \Box\Diamond^{n+1}\neg\varphi \vdash \Box\varphi \wedge \Diamond^{n+1}\neg\varphi \vdash \Box\varphi \wedge \Diamond^n\neg\Box\varphi \vdash \bot,$$

using the inductive hypothesis for the last step. Thus, $\varphi \wedge \Diamond^{n+1}\neg\varphi \vdash \bot$ by Lemma 3.25. □

3.2.4 Distinguishing Boolean propositions

We now make precise the idea from § 2.5 that classical principles should hold for the "non-modal fragment." If we understand the basic elements $p, q, r \in \mathsf{Prop}$ of our inductively defined formal language as *propositional variables*, standing in for arbitrary propositions including epistemic modal propositions, then we do not want classical principles like $p \wedge (q \vee r) \vdash (p \wedge q) \vee (p \wedge r)$, since accepting such a principle for p, q, r means

[12] A potential confound in these cases comes from the phenomenon of modal concord, where different modals do not contribute independent modal meanings; see van Wijnbergen-Huitink 2020 for an overview. Moss's claim is that these cases are not cases of modal concord; judgments are, to be sure, somewhat tenuous.

accepting it for all propositions (cf. Burgess 2003, pp. 147-8). However, we can add to the basic elements of our formal language a set Bool whose elements we think of as variables only for *non-epistemic*, or *Boolean*, propositions.[13] Typographically, $p, q, r, \ldots$ are elements of Prop, whereas $\mathtt{p}, \mathtt{q}, \mathtt{r}, \ldots$ are elements of Bool.

Definition 3.29. Let $\mathcal{EL}^+$ be the language generated by the grammar

$$\varphi ::= \top \mid p \mid \mathtt{p} \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \Box\varphi$$

where p belongs to a countably infinite set Prop and $\mathtt{p}$ belongs to a countably infinite set Bool. A formula of $\mathcal{EL}^+$ is said to be *Boolean* if all its propositional variables are from Bool and it does not contain $\Box$.

In line with our idea that the only failures of classicality come from epistemic modals, we interpret the variables of Bool in a Boolean subalgebra of our ambient ortholattice of propositions.

Definition 3.30. A *modal ortho-Boolean lattice* is a tuple $\langle A, B, \vee, 0, \wedge, 1, \neg, \Box\rangle$ where $\langle A, \vee, 0, \wedge, 1, \neg, \Box\rangle$ is a modal ortholattice and $\langle B, \vee_{|B}, 0, \wedge_{|B}, 1, \neg_{|B}\rangle$ is a Boolean algebra where $B \subseteq A$ and $\vee_{|B}$, $\wedge_{|B}$, and $\neg_{|B}$ are the restrictions of $\vee$, $\wedge$, and $\neg$, respectively, to B.

Note that every modal ortholattice can be expanded to a modal ortho-Boolean lattice by taking $B = \{0, 1\}$, so modal ortho-Boolean lattices may be viewed as a generalization of modal ortholattices.

So far there is no required connection between the distinguished Boolean subalgebra of a modal ortho-Boolean lattice and the ambient modal ortholattice: any Boolean subalgebra can be distinguished. However, there appear to be intuitively valid principles relating Boolean—but not arbitrary—propositions and epistemic modals, and capturing such principles requires a certain coherence between the Boolean subalgebra B and the ambient epistemic ortholattice. Consider, for example, the following.

Definition 3.31. Given a modal ortho-Boolean lattice $L = \langle A, B, \vee, 0, \wedge, 1, \neg, \Box\rangle$, define:

- $B_0 = B$;
- B_{n+1} is the subortholattice of $\langle A, \vee, 0, \wedge, 1, \neg\rangle$ generated by $\{\Box b \mid b \in B_n\}$.

Then L is *level-wise Boolean* if B_n is Boolean for each $n \in \mathbb{N}$.

The motivation for this condition is straightforward: no natural language counterexample to a classical inference that we have found is such that all propositions come from the same level B_n. For example, the counterexample to pseudocomplementation, going from $p \wedge \Diamond\neg p = 0$ to $\Diamond\neg p \leq \neg p$, involves $p, \neg p \in B_n$ and $\Diamond\neg p \in B_{n+1}$. Similar points apply to our counterexamples involving distributivity, disjunctive syllogism, and orthomodularity. This observation is, to our knowledge, novel, and suggests a picture on which classical reasoning *across different epistemic levels* can be dangerous, but classical reasoning within a given epistemic level is safe. We can model that picture with level-wise Boolean epistemic ortholattices:

Definition 3.32. An *epistemic ortho-Boolean lattice* is a modal ortho-Boolean lattice $\langle A, B, \vee, 0, \wedge, 1, \neg, \Box\rangle$ that is level-wise Boolean and such that $\langle A, \vee, 0, \wedge, 1, \neg, \Box\rangle$ is an epistemic ortholattice.

For checking the level-wise Boolean condition, note that if $B_{n+1} = B_n$, then $B_k = B_n$ for all $k \geq n$; and if A is finite, then we are bound to reach such a fixed point B_n.

[13] The distinction between viewing atomic formulas of the propositional modal language as genuine propositional variables vs. non-epistemic variables has a precedent in the literature on dynamic epistemic logic (see, e.g., Holliday et al. 2012, 2013).

Example 3.33. Consider the epistemic ortholattice in Figure 3 with $B = \{\bot, p, \neg p, \top\}$, highlighted on the left of Figure 4. Then B forms a subortholattice and a four-element Boolean algebra. Moreover, $B_1 = \{\bot, \Box p, \Diamond\neg p, \Diamond p, \Box\neg p, \Box p \vee \Box\neg p, \Diamond p \wedge \Diamond\neg p, \top\}$ forms an eight-element Boolean algebra, highlighted on the right of Figure 4; and $B_2 = B_1$, so $B_n = B_1$ for all $n \geq 1$. Thus, by equipping the epistemic ortholattice in Figure 3 with B, we obtain an epistemic ortho-Boolean lattice.

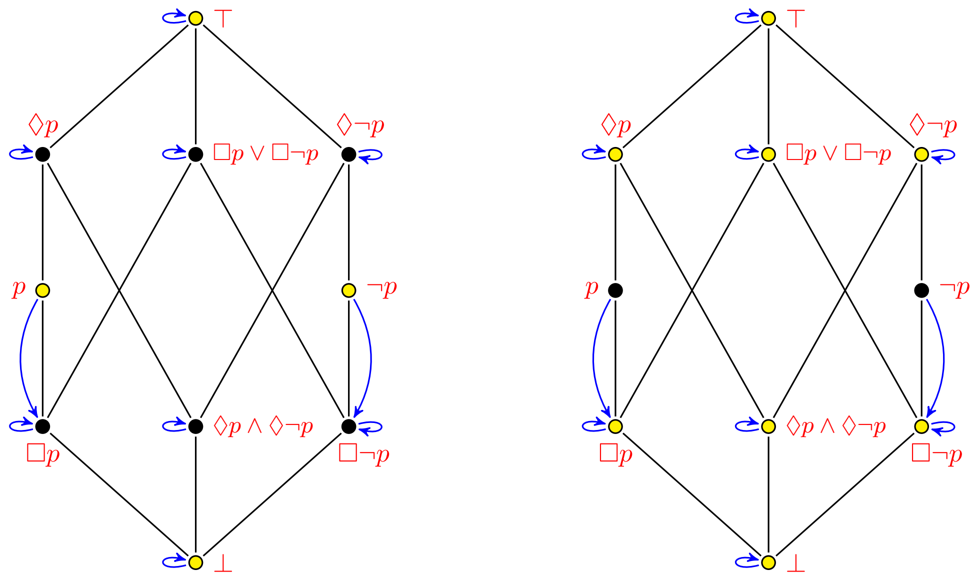

Figure 4: Highlighting in yellow the Boolean subalgebras B_0 (left) and B_1 (right) of the epistemic ortholattice from Figure 3.

Let us now turn to the semantics of $\mathcal{EL}^+$.

Definition 3.34. A valuation θ on a modal ortho-Boolean lattice $\langle A, B, \vee, 0, \wedge, 1, \neg, \Box\rangle$ is a map $\theta : \mathsf{Prop} \cup \mathsf{Bool} \to A$ such that for all $\mathsf{p} \in \mathsf{Bool}$, $\theta(\mathsf{p}) \in B$. Such a valuation extends to $\tilde{\theta} : \mathcal{EL}^+ \to A$ by: $\tilde{\theta}(\top) = 1$, $\tilde{\theta}(\neg\varphi) = \neg\tilde{\theta}(\varphi)$, $\tilde{\theta}(\varphi \wedge \psi) = \tilde{\theta}(\varphi) \wedge \tilde{\theta}(\psi)$, and $\tilde{\theta}(\Box\varphi) = \Box\tilde{\theta}(\varphi)$.

Corresponding to the different subortholattices B_n, we have a hierarchy of sublanguages $\mathcal{B}_n$.

Definition 3.35.

- Let $\mathcal{B}_0$ be the set of Boolean formulas as in Definition 3.29.
- Let $\mathcal{B}_{n+1}$ be the smallest set of formulas that includes $\{\Box\varphi \mid \varphi \in \mathcal{B}_n\}$ and is closed under $\neg$ and $\wedge$.

An obvious induction on the structure of formulas shows that formulas at each level are interpreted in the corresponding B_n.

Lemma 3.36. Let θ be a valuation on a modal ortho-Boolean lattice $\langle A, B, \vee, 0, \wedge, 1, \neg, \Box\rangle$. Then for any $n \in \mathbb{N}$ and $\beta \in \mathcal{B}_n$, we have $\tilde{\theta}(\beta) \in B_n$.

Our consequence relation for $\mathcal{EL}^+$ is defined just like our consequence relation for $\mathcal{EL}$ but quantifying over modal ortho-Boolean lattices instead of modal ortholattices.

Definition 3.37. Given a class $\mathbf{C}$ of modal ortho-Boolean lattices, define the semantic consequence relation $\vDash^+_{\mathbf{C}}$, a binary relation on $\mathcal{EL}^+$, as follows: $\varphi \vDash^+_{\mathbf{C}} \psi$ if for every $L \in \mathbf{C}$ and valuation θ on L, we have $\tilde{\theta}(\varphi) \leq \tilde{\theta}(\psi)$, where $\leq$ is the lattice order of L.

Turning to the logic, we now explicitly include the principle of distributivity for formulas at a given epistemic level.[14]

Definition 3.38. Let EO^+ be defined just like EO in Definition 3.24 for arbitrary formulas $\varphi, \psi, \chi \in \mathcal{EL}^+$ but in addition we have for all $n \in \mathbb{N}$ and $\alpha, \beta, \gamma \in \mathcal{B}_n$:

16. $\alpha \wedge (\beta \vee \gamma) \vdash (\alpha \wedge \beta) \vee (\alpha \wedge \gamma)$.

Adding distributivity for formulas at a given epistemic level allows us to derive all principles of classical logic for formulas at that level. For example, we can derive the psueodcomplementation principle that $\alpha \wedge \beta \vdash \bot$ implies $\alpha \vdash \neg\beta$ for $\alpha, \beta \in \mathcal{B}_n$ as in the proof of Proposition 3.7.

We can also derive principles of classical normal modal logic for formulas at a given epistemic level.

Example 3.39. We can derive that $\Diamond\delta \wedge \Box\gamma \vdash \Diamond(\delta \wedge \gamma)$ for $\delta, \gamma \in \mathcal{B}_n$ as follows:

$$\Diamond\delta \wedge \Box\gamma \wedge \Box\neg(\delta \wedge \gamma) \vdash \Diamond\delta \wedge \Box(\gamma \wedge \neg(\delta \wedge \gamma)) \vdash \Diamond\delta \wedge \Box\neg\delta \vdash \bot,$$

so by pseudocomplementation for formulas in $\mathcal{B}_{n+1}$, we have $\Diamond\delta \wedge \Box\gamma \vdash \Diamond(\delta \wedge \gamma)$. Note that we would not want this principle for formulas of different epistemic levels, since $\Diamond p \wedge \Box\Diamond\neg p$ should not entail $\Diamond(p \wedge \Diamond\neg p)$.

More generally, we have the following completeness theorem.

Theorem 3.40. The logic EO^+ (summarized in Figure 5) is sound and complete with respect to the class $\mathbf{EO}^+$ of all epistemic ortho-Boolean lattices according to the consequence relation of Definition 3.37: for all $\varphi, \psi \in \mathcal{EL}^+$, we have $\varphi \vdash_{\mathsf{EO}^+} \psi$ if and only if $\varphi \vDash^+_{\mathbf{EO}} \psi$.

Proof. Soundness is again straightforward. For completeness, as in the proofs of Theorems 3.13 and 3.26, we consider the Lindenbaum-Tarski algebra L of EO^+. Let $B = \{[\beta] \mid \beta \text{ Boolean}\}$, and note that B forms a subortholattice of L. As a consequence of the distributivity rule of EO^+, each B_n generated from B is a Boolean algebra under the restricted operations of L. Hence we have an epistemic ortho-Boolean lattice. Finally, let θ be the valuation on L with $\theta(p) = [p]$ for all $p \in \mathsf{Prop}$ and $\theta(\mathsf{p}) = [\mathsf{p}]$ for all $\mathsf{p} \in \mathsf{Bool}$. The rest of the proof is the same as for Theorems 3.13 and 3.26. ☐

As before, soundness and completeness with respect to $\mathsf{S5}$ epistemic ortho-Boolean lattices is also straightforward by adding the rules $\Box\varphi \vdash \Box\Box\varphi$ and $\Diamond\varphi \vdash \Box\Diamond\varphi$ to Definition 3.38.

Thus, EO^+ is our proposed base logic for epistemic modals, which satisfies all the desiderata of § 2. We summarize the principles of EO^+ in Figure 5. In a sense, we could declare "mission accomplished" now. However, while the algebraic semantics shows perspicuously exactly how we are departing from the Boolean algebraic semantics underlying classical logic, some may be unsatisfied with the way in which the algebraic approach simply builds into our semantics—in the form of equations or inequalities on lattices—precisely the logical principles we want to get out. Thus in the next section, we give a more concrete semantics, defining a notion of *truth at a possibility*, that delivers the same logic without explicitly building logical principles into the semantics.

[14] A Fitch-style proof system for EO^+ can be obtained from the Fitch-style proof system for EO sketched in Footnote 11 by adding a rule of Level-wise Reiteration, which is like Fitch's rule of Reiteration except that it only allows reiteration of formulas from $\mathcal{B}_n$ into subproofs whose opening assumptions are from $\mathcal{B}_n$.

1. $\varphi \vdash \top$;
2. $\varphi \vdash \varphi$;
3. $\varphi \wedge \psi \vdash \varphi$;
4. $\varphi \wedge \psi \vdash \psi$;
5. $\varphi \vdash \neg\neg\varphi$;
6. $\neg\neg\varphi \vdash \varphi$;
7. $\varphi \wedge \neg\varphi \vdash \psi$;
8. if $\varphi \vdash \psi$ and $\psi \vdash \chi$, then $\varphi \vdash \chi$;
9. if $\varphi \vdash \psi$ and $\varphi \vdash \chi$, then $\varphi \vdash \psi \wedge \chi$;
10. if $\varphi \vdash \psi$, then $\neg\psi \vdash \neg\varphi$.

11. if $\varphi \vdash \psi$, then $\Box\varphi \vdash \Box\psi$;
12. $\Box\varphi \wedge \Box\psi \vdash \Box(\varphi \wedge \psi)$;
13. $\varphi \vdash \Box\top$;
14. $\Box\varphi \vdash \varphi$;
15. $\neg\varphi \wedge \Diamond\varphi \vdash \bot$;

16. $\alpha \wedge (\beta \vee \gamma) \vdash (\alpha \wedge \beta) \vee (\alpha \wedge \gamma)$ for $\alpha, \beta, \gamma \in \mathcal{B}_n$.

Figure 5: Principles of the epistemic orthologic EO^+

4 Possibility semantics

In this section, we turn from abstract algebraic semantics to a more concrete and perhaps more intuitive and illuminating *possibility semantics* for epistemic orthologic. We begin with possibility semantics for non-modal orthologic, reviewing the approach of Goldblatt 1974.[15] Although Goldblatt does not classify his semantics under the banner of "possibility semantics," this classification is justified by the close connection, explained in Holliday 2021a, between his semantics for orthologic using a relation of *compatibility* (or what he calls *proximity*) and possibility semantics for classical logic using a relation of *refinement*.[16] Indeed, classical possibility semantics as developed in Humberstone 1981 and Holliday 2014, Forthcoming, 2021b can be recast in terms of compatibility instead of refinement (see Remark 4.9 and Holliday 2022). Here we assume no previous exposure to possibility semantics of any kind. All we assume is familiarity with possible world semantics, of which possibility semantics is a generalization.

4.1 Review of possibility semantics for orthologic

Possibility semantics for orthologic replaces the set W of worlds from classical possible world semantics with a set S of possibilities, endowed with a *compatibility relation* $\between$ between possibilities.

Definition 4.1. A (*symmetric*) *compatibility frame* is a pair $\mathcal{F} = \langle S, \between \rangle$ where S is a nonempty set and $\between$ is a reflexive and symmetric binary relation on S.

A set W of possible worlds may be regarded as a compatibility frame in which each $w \in W$ is compatible only with itself: $v \between w$ implies $v = w$. Think of $x \between y$ as meaning that x *does not settle as true anything that* y *settles as false*. If w and v are complete possible worlds, then—assuming distinct worlds must differ in some respect—w must settle something as true that y settles as false. However, unlike possible worlds, distinct *possibilities* may be compatible with each other—a sign of their *partiality*. For example, the possibility that it is raining in Beijing is compatible with the possibility that it is sunny in Malibu.

[15] A similar approach appears in Dishkant 1972; for comparison of the two approaches, see Holliday 2021a, Remark 2.12.

[16] In fact, Goldblatt focuses on an equivalent semantics for orthologic using the complement of compatibility, namely *incompatibility* or *orthogonality*, but whether one works with compatibility or incompatibility is just a matter of preference.

The above notion of compatibility is weaker than what might be called *compossibility*: x and y are compossible if there is some possibility z that settles as true everything that x settles as true and everything that y settles as true. As we shall see below, in the above sense of compatibility, a possibility x settling as true that *it might be raining in Beijing* can be *compatible* with a possibility y settling as true that *it is not raining in Beijing*, since settling as true *that it might be raining in Beijing* does not entail settling as false *that it is not raining in Beijing*; however, such an x and y cannot be *compossible*, since there can be no possibility z that settles as true that *it might be raining in Beijing but it isn't*.

Remark 4.2. Holliday 2021a, 2022 considers more general compatibility frames in which $\between$ is only assumed to be reflexive,[17] as such frames can be used to represent arbitrary complete lattices; but here we use the term 'compatibility frame' for only the symmetric frames, which will give rise to complete ortholattices.

We now turn to the question of what counts as a *proposition*. In basic possible world semantics for classical logic, every set of worlds is (or corresponds to) a proposition; as a result, the collection of propositions ordered by inclusion forms a Boolean algebra. In semantics for *non-classical* logics, not every set of worlds or possibilities can count as a proposition—for then we would still be stuck with a Boolean algebra of propositions. Thus, for example, in possible world semantics for intuitionistic logic (Dummett and Lemmon 1959, Grzegorczyk 1964, Kripke 1965), only sets of worlds that are upward closed with respect to an information order count as legitimate propositions. In possibility semantics for orthologic, only sets of possibilities that behave sensibly with respect to the compatibility relation count as legitimate propositions.

Definition 4.3. Given a compatibility frame $\langle S, \between\rangle$, a set $A \subseteq S$ is said to be $\between$*-regular* if for all $x \in S$,

$$x \notin A \Rightarrow \exists y \between x \; \forall z \between y \;\; z \notin A.$$

The idea is that if x does *not* make a proposition A true, then there should be a possibility y compatible with x that makes A *false*, so that all possibilities z compatible with y do not make A true.[18] See Figure 6. If x already makes A false, so no possibility compatible with x makes A true, then we can simply take $y = x$. The interesting case is where x neither makes A true nor makes A false. Thus, regularity can be expressed in the slogan:

Indeterminacy Implies Compatibility with Falsity.

Thus, if A is indeterminate at x, then x is compatible with a y that makes A false.

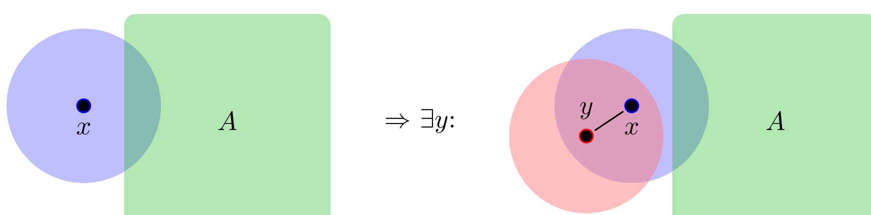

Figure 6: A depiction of the $\between$-regularity of A: if x is not in A, then x is compatible with a y that is not compatible with any possibility in A. The blue disk represents the set of possibilities compatible with x; the red disk represents the set of possibilities compatible with y; and the green region represents A.

[17] In this case, Definition 4.3 below is modified to use '$\forall z \between^{-1} y$' instead of '$\forall z \between y$'.

[18] Another way to view the definition of $\between$-regular sets is via the function $c_{\between} : \wp(S) \to \wp(S)$ defined as follows: for $A \subseteq S$, we have $c_{\between}(A) = \{x \in S \mid \forall y \between x \; \exists z \between y : z \in A\}$. This $c_{\between}$ is a *closure operator* (i.e., inflationary, idempotent, and monotone operation) on $\wp(S)$. The *fixpoints* of $c_{\between}$, i.e., those A such that $A = c_{\between}(A)$, are exactly the $\between$-regular sets as defined above.

From the compatibility relation, it is useful to define a relation of *refinement*. This can be done in two equivalent ways: y refines x if every proposition settled true by x is also settled true by y, or equivalently, if every possibility compatible with y is compatible with x.

Lemma 4.4. For any compatibility frame $\langle S, \between \rangle$, the following are equivalent for any $x, y \in S$:

1. for all $\between$-regular sets $A \subseteq S$, if $x \in A$, then $y \in A$;
2. for all $z \in S$, if $z \between y$ then $z \between x$.

When these conditions hold, we write $y \sqsubseteq x$.

Proof. From 1 to 2, suppose $z \between y$ but not $z \between x$. Let $A = \{v \in S \mid \text{not } z \between v\}$. We claim that A is $\between$-regular. We must show that if $u \notin A$, then $\exists u' \between u\ \forall u'' \between u'\ u'' \notin A$, i.e., that if $z \between u$, then $\exists u' \between u\ \forall u'' \between u'\ z \between u''$. Indeed, let $u' = z$ and use the symmetry of $\between$. Thus, A is $\between$-regular, and $x \in A$ but $y \notin A$.

From 2 to 1, let A be a $\between$-regular set. Suppose $x \in A$ and condition 2 holds. We claim that for all $y' \between y$ there is a $y'' \between y'$ with $y'' \in A$, so $y \in A$ by the $\between$-regularity of A. Suppose $y' \between y$. Then since condition 2 holds, we have $y' \between x$ and hence $x \between y'$ by symmetry of $\between$. Thus, we may take $y'' = x$. □

For any possibility x, the set of refinements of x is a $\between$-regular set and hence may be regarded as a proposition—the proposition *that possibility x obtains*.

Lemma 4.5. Given a compatibility frame $\mathcal{F} = \langle S, \between \rangle$ and $x \in S$, the set $\downarrow x = \{y \in S \mid y \sqsubseteq x\}$ is $\between$-regular.

Proof. Suppose $y \notin \downarrow x$. Hence there is a $y' \between y$ such that *not* $y' \between x$. Now consider any $y'' \between y'$. Since $y' \between y''$ but *not* $y' \between x$, we have that $y'' \not\sqsubseteq x$. Thus, we have shown that if $y \notin \downarrow x$, then $\exists y' \between y\ \forall y'' \between y'\ y'' \notin \downarrow x$. Hence $\downarrow x$ is $\between$-regular. □

We can also use the notion of refinement to define *possible worlds*: a world is a possibility that refines every possibility with which it is compatible.

Definition 4.6. Given a compatibility frame $\mathcal{F} = \langle S, \between \rangle$, a *world in $\mathcal{F}$* is a $w \in S$ such that for all $x \in S$, $w \between x$ implies $w \sqsubseteq x$.

Now in line with the idea that only $\between$-regular sets of possibilities are propositions, a *model* should interpret the propositional variables of our formal language as $\between$-regular sets.

Definition 4.7. A *compatibility model* is a pair $\mathcal{M} = \langle \mathcal{F}, V \rangle$ where $\mathcal{F} = \langle S, \between \rangle$ is a compatibility frame and V is a function assigning to each $p \in \mathsf{Prop}$ a $\between$-regular set $V(p) \subseteq S$. We say that $\mathcal{M}$ is *based on* $\mathcal{F}$.

We now observe how any compatibility frame gives rise to an ortholattice. Compare the following result, due to Birkhoff (1940, §§ 32-4), to the way in which for any set W, the family of all subsets ordered by the subset relation $\subseteq$ forms a complete Boolean algebra.

Proposition 4.8. Given any compatibility frame $\mathcal{F} = \langle S, \between \rangle$, the $\between$-regular sets ordered by the subset relation $\subseteq$ form a complete lattice, which becomes an ortholattice with the operation $\neg$ defined by

$$\neg A = \{x \in S \mid \forall y \between x\ \ y \notin A\}. \tag{I}$$

For the lattice operations, we have $A \wedge B = A \cap B$, $A \vee B = \neg(\neg A \cap \neg B)$, $1 = S$, and $0 = \varnothing$. We denote the ortholattice of $\between$-regular subsets of $\mathcal{F}$ by $O(\mathcal{F})$.

Thus, $\wedge$ is intersection as in possible world semantics, but note how $\neg$ and $\vee$ are no longer interpreted as in possible world semantics. In possible world semantics, $\neg A = \{x \in W \mid x \notin A\}$, in contrast to (I), and $\vee$ is union, in contrast to what we get when we unpack the definition $\vee$ as $A \vee B = \neg(\neg A \cap \neg B)$:

$$A \vee B = \{x \in S \mid \forall y \between x \, \exists z \between y : z \in A \cup B\}, \tag{II}$$

i.e., x makes $A \vee B$ true just in case every possibility compatible with x is in turn compatible with a possibility that makes one of A or B true.

Remark 4.9. As observed in Holliday 2021a (Example 3.16), the ortholattice arising from a compatibility frame as in Proposition 4.8 is a Boolean algebra if for all $x, y \in S$, if $x \between y$, then x and y have a common refinement, i.e., there is a z such that $z \sqsubseteq x$ and $z \sqsubseteq y$. Indeed, this condition is not only sufficient but also necessary for the ortholattice to be Boolean: for if x and y have no common refinement, so $\downarrow x \cap \downarrow y = \varnothing$, then in a Boolean algebra we must have $\downarrow x \subseteq \neg \downarrow y$, contradicting $x \between y$. Thus, classicality corresponds to the condition (to use terminology introduced after Definition 4.1) that *compatibility* implies *compossibility*.

Our departure from classicality can now be seen as follows: we want there to be distinct but *compatible* possibilities that settle A and $\Diamond\neg A$ as true, respectively, as neither A nor $\Diamond\neg A$ should entail the negation of the other; yet such possibilities should not be *compossible*—they should have no common refinement, since no single possibility should settle both A and $\Diamond\neg A$ as true.

Remark 4.10. It is noteworthy that two compatibility frames $(S, \between)$ and $(S, \between')$ can have the same derived refinement relation but give rise to non-isomorphic ortholattices; an example is provided in the second of the notebooks cited in § 1. Thus, although the refinement relation has all the information one needs in a classical setting (wherein compatibility can be defined from a primitive partial order of refinement by: $x \between y$ if x and y have a common refinement), it does not in our non-classical setting here.

Example 4.11. Figure 7 shows a simple compatibility frame with five possibilities (above) and its refinement relation (below) defined from compatibility as in Lemma 4.4.

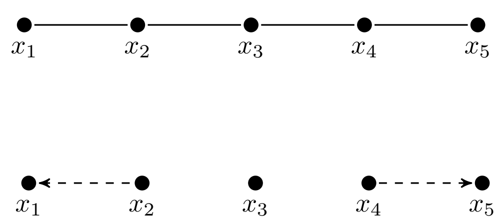

Figure 7: A compatibility frame (above), where compatible possibilities are linked by an edge in the graph, and its refinement relation (below), where a dashed arrow from y to z indicates that z is a refinement of y, i.e., $z \sqsubseteq y$. All reflexive compatibility loops and reflexive refinement loops are omitted.

Figure 8 shows the ten $\between$-regular subsets of the compatibility frame from Figure 7, each highlighted in green. For example, $\{x_1, x_2, x_3\}$ is $\between$-regular according to Definition 4.3 because each possibility outside the set is compatible with x_5, which is not compatible with anything in the set; but $\{x_2, x_3, x_4\}$ is not $\between$-regular, because x_1 is outside the set and yet everything compatible with x_1 is compatible with something in the set. Note that ordering the $\between$-regular subsets in Figure 8 by $\subseteq$ and defining $\neg$ as in Proposition 4.8 yields an ortholattice $O(\mathcal{F})$ isomorphic to the ortholattice in Figure 3 (reproduced on the right of Figure 10 below)!

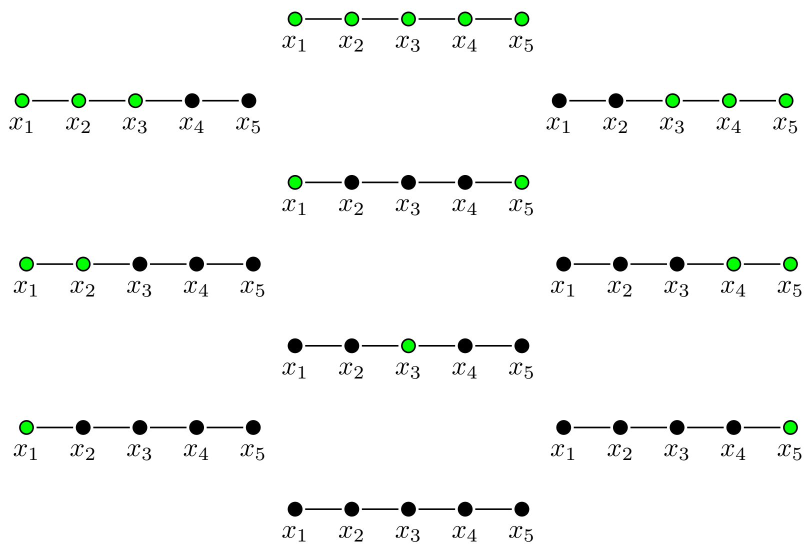

Figure 8: The ten $\between$-regular subsets of the compatibility frame from Figure 7, each highlighted in green

Example 4.12. The two ortholattices in Figure 1 can be realized as the ortholattices of $\between$-regular subsets of the two compatibility frames in Figure 9. Verifying this claim is a good check for understanding.

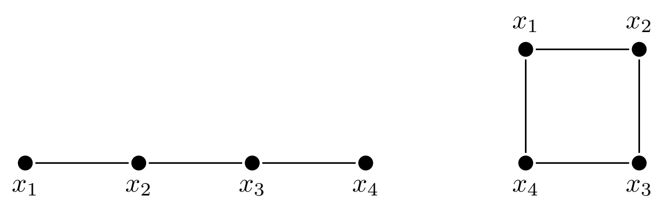

Figure 9: Compatibility frames that give rise to the ortholattices in Figure 1 (reproduced on the left and middle of Figure 10)

From Examples 4.11 and 4.12, one might notice the key to representing any finite ortholattice using a compatibility frame: the possibilities in the frame correspond to the *join-irreducible* elements of the ortholattices, i.e., those nonzero elements a of the ortholattice that cannot be obtained as a join of elements distinct from a—intuitively, those noncontradictory propositions that cannot be expressed as a disjunction of distinct propositions; then two possibilities a and b are compatible if $a \not\leq \neg b$ in the ortholattice. Figure 10 highlights in orange the join-irreducible elements of the three ortholattices considered so far, elements which one can match up with the possibilities in the corresponding compatibility frames in Figures 9 and 7.

The fact that finite ortholattices can be represented using compatibility frames as described above is a consequence of the following more general representation theorem for all complete ortholattices.[19] A set V of elements of a lattice L is said to be *join-dense in* L if every element of L can be obtained as the (possibly infinite) join of some elements of V (e.g., the set of all elements of L is trivially join-dense in L).

[19]For the representation of arbitrary (including incomplete) ortholattices using compatibility (or incompatibility) frames equipped with a topology, see Goldblatt 1975 and McDonald and Yamamoto 2021, and for associated categorical dualities, see Bimbó 2007, Dmitrieva 2021, and McDonald and Yamamoto 2021.

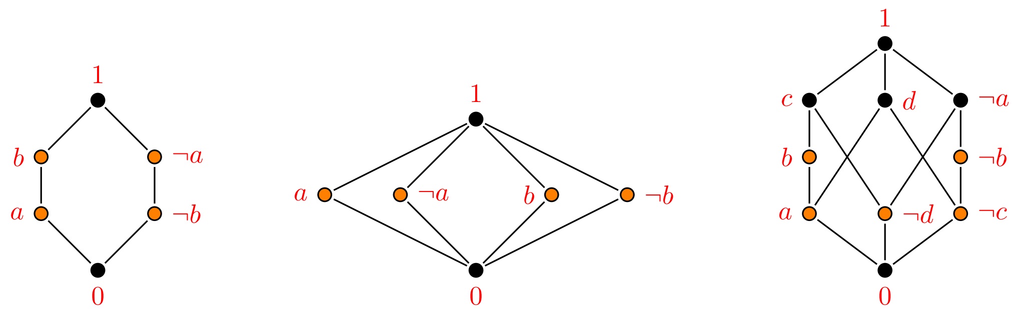

Figure 10: Hasse diagrams of ortholattices from Figures 1 and 3 with join-irreducible elements in orange

Theorem 4.13. Let $L = \langle A, \vee, 0, \wedge, 1, \neg \rangle$ be an ortholattice and V any join-dense subset of L. Then where $\mathcal{F} = \langle V \setminus \{0\}, \between \rangle$ is the compatibility frame based on $V \setminus \{0\}$ with $\between$ defined by $a \between b$ if $a \nleq \neg b$, we have that L embeds into $O(\mathcal{F})$ via the map $a \mapsto \{b \in V \setminus \{0\} \mid b \leq a\}$, which is an isomorphism if L is complete.

For a proof, see, e.g., MacLaren 1964, Theorems 2.3 and 2.5. Compare this representation of complete ortholattices to Tarski's (1935) representation of *complete and atomic Boolean algebras*, which lies at the foundation of possible world semantics: any such Boolean algebra is isomorphic to the powerset of its set of atoms (think possible worlds).[20] Now in a *finite* ortholattice, the set of join-irreducible elements is join-dense in the lattice, so we obtain the following corollary of Theorem 4.13, which explains Examples 4.11 and 4.12.

Corollary 4.14. Let $L = \langle A, \vee, 0, \wedge, 1, \neg \rangle$ be a finite ortholattice. Then where $\mathcal{F} = \langle J, \between \rangle$ is the compatibility frame based on the set J of join-irreducible elements of L with $\between$ defined by $a \between b$ if $a \nleq \neg b$, we have that $O(\mathcal{F})$ is isomorphic to L.

Let us now turn to formal semantics for the language $\mathcal{L}$ (recall Definition 3.9). Proposition 4.8 leads us to the following.

Definition 4.15. Given a compatibility model $\mathcal{M} = \langle S, \between, V \rangle$, $x \in S$, and $\varphi \in \mathcal{L}$, we define $\mathcal{M}, x \Vdash \varphi$ as follows:

1. $\mathcal{M}, x \Vdash \top$;
2. $\mathcal{M}, x \Vdash p$ iff $x \in V(p)$;
3. $\mathcal{M}, x \Vdash \varphi \wedge \psi$ iff $\mathcal{M}, x \Vdash \varphi$ and $\mathcal{M}, x \Vdash \psi$;
4. $\mathcal{M}, x \Vdash \neg\varphi$ iff for all $y \between x$, $\mathcal{M}, y \nVdash \varphi$.

We define $[\![\varphi]\!]^{\mathcal{M}} = \{x \in S \mid \mathcal{M}, x \Vdash \varphi\}$.

Then given our definition of $\varphi \vee \psi$ as $\varphi \vee \psi := \neg(\neg\varphi \wedge \neg\psi)$, as in (II) we have:

- $\mathcal{M}, x \Vdash \varphi \vee \psi$ iff for all $y \between x$ there is a $z \between y$ such that $\mathcal{M}, z \Vdash \varphi$ or $\mathcal{M}, z \Vdash \psi$.

An easy induction shows that the set of possibilities that make a formula true is indeed a proposition.

Lemma 4.16. For any compatibility model $\mathcal{M}$ and $\varphi \in \mathcal{L}$, $[\![\varphi]\!]^{\mathcal{M}}$ is a $\between$-regular set.

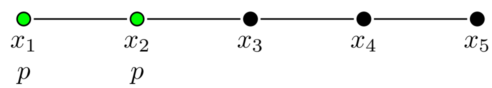

Figure 11: A compatibility model based on the frame from Figure 7 with $V(p)$ highlighted in green

Example 4.17. Consider a compatibility model based on the frame in Figure 7 with $V(p) = \{x_1, x_2\}$, so $V(p)$ is a $\between$-regular set, as shown in Figure 11. Then observe that $[\![\neg p]\!]^{\mathcal{M}} = \{x_4, x_5\}$, another $\between$-regular set.

We define semantic consequence as usual in terms of truth preservation.

Definition 4.18. Given a class $\mathbf{C}$ of compatibility frames, define the semantic consequence relation $\vDash_{\mathbf{C}}$, a binary relation on $\mathcal{L}$, as follows: $\varphi \vDash_{\mathbf{C}} \psi$ if for every $\mathcal{F} \in \mathbf{C}$, model $\mathcal{M}$ based on $\mathcal{F}$, and possibility x in $\mathcal{M}$, if $\mathcal{M}, x \Vdash \varphi$, then $\mathcal{M}, x \Vdash \psi$.

We can of course also define a consequence relation between a set of premises on the left and a single conclusion on the right: $\Gamma \vDash_{\mathbf{C}} \psi$ if for every $\mathcal{F} \in \mathbf{C}$, model $\mathcal{M}$ based on $\mathcal{F}$, and possibility x in $\mathcal{M}$, if $\mathcal{M}, x \Vdash \varphi$ for all $\varphi \in \Gamma$, then $\mathcal{M}, x \Vdash \psi$. But for simplicity we only consider finite sets of premises here, in which case a single premise on the left suffices given that $\mathcal{L}$ contains a conjunction interpreted as intersection.

Goldblatt (1974) proved the following completeness theorem, showing that the possibility semantics above is indeed a semantics for the orthologic O of Definition 3.12.

Theorem 4.19. The logic O is sound and complete with respect to the class $\mathbf{CF}$ of all compatibility frames according to the consequence relation in Definition 4.18: for all $\varphi, \psi \in \mathcal{L}$, we have $\varphi \vdash_{\mathsf{O}} \psi$ if and only if $\varphi \vDash_{\mathbf{CF}} \psi$.

Theorem 4.19 is easily proved from the embedding result in Theorem 4.13. In § 4.2 we will give a proof of a modal version of this completeness theorem that will show another way to prove Theorem 4.19.

4.2 Adding epistemic modality

In this section, we extend the compatibility frames of § 4.1 to interpret modalities $\Box$ and $\Diamond$. We do so in the style of *relational* possibility semantics (Humberstone 1981, Holliday Forthcoming, 2021b, 2022); as we note below, we could equally work with a *functional possibility semantics* (Holliday 2014) instead.

As usual, given a binary relation R on a set S of possibilities, we define a $\Box$ operation on propositions by

$$\Box A = \{x \in S \mid R(x) \subseteq A\}, \tag{III}$$

where $R(x) = \{y \in S \mid xRy\}$. Given our definition of $\Diamond A$ as $\neg\Box\neg A$, we have

$$\Diamond A = \{x \in S \mid \forall x' \between x \; \exists y' \in R(x') \; \exists y'' \between y' : y'' \in A\}. \tag{IV}$$

Finally, we posit one condition on the relation between *epistemic accessibility* and what *might* be the case:

- R-regularity: if x can epistemically access a possibility compatible with y, then x is compatible with a possibility according to which y *might obtain*.

[20]Recall that an *atom* in a bounded lattice is a nonzero element a such that $b \leq a$ implies $b = 0$ or $b = a$. A lattice L is *atomic* if for every nonzero element b in L, there is an atom $a \leq b$.

Formally, the antecedent means that $xRy' \between y$, and the consequent means that x is compatible with a possibility $x' \in \Diamond{\downarrow}y$ (recall Lemma 4.5).

We can now define the basic frames for our modal semantics.

Definition 4.20. A *modal compatibility frame* is a triple $\mathcal{F} = \langle S, \between, i\rangle$ where $\langle S, \between\rangle$ is a compatibility frame and R is a binary relation on S satisfying

- R-regularity: if $xRy' \between y$, then $\exists x' \between x$: $x' \in \Diamond{\downarrow}y$.

A *modal compatibility model* is a pair $\mathcal{M} = \langle \mathcal{F}, V\rangle$ where $\mathcal{F} = \langle S, \between, i\rangle$ is a modal compatibility frame and $\langle S, \between, V\rangle$ is a compatibility model as in Definition 4.7. We say that $\mathcal{M}$ is *based on* $\mathcal{F}$.

Lemma 4.21. R-regularity can be written equivalently without $\Diamond$ as follows:

- R-regularity: if $xRy' \between y$, then $\exists x' \between x\ \forall x'' \between x'\ \exists y''$: $x''Ry'' \between y$.

Proof. By (IV), $x' \in \Diamond{\downarrow}y$ in R-regularity is equivalent to: $\forall x'' \between x'\ \exists y'' \in R(x'')\ \exists y''' \between y''$: $y''' \in {\downarrow}y$. But $\exists y''' \between y''$: $y''' \in {\downarrow}y$ is equivalent to $y'' \between y$; from left-to-right, if $y''' \sqsubseteq y$, then $y'' \between y'''$ implies $y'' \between y$, and from right-to-left, take $y''' = y$. Thus, $x' \in \Diamond{\downarrow}y$ is equivalent to the condition on x' in the lemma. ☐

We will begin by proving using R-regularity that if A is a proposition, so is $\Box A$.

Proposition 4.22. For any modal compatibility frame $\mathcal{F} = \langle S, \between, i\rangle$ and $\between$-regular set $A \subseteq S$, we have that $\Box A$ is $\between$-regular.

Proof. We must show that

$$x \notin \Box A \Rightarrow \exists x' \between x\, \forall x'' \between x'\ \ x'' \notin \Box A.$$

Suppose $x \notin \Box A$, so for some $y \in R(x)$, we have $y \notin A$. Then since A is $\between$-regular, there is a $z \between y$ such that $z \in \neg A$. Since $xRy \between z$, by R-regularity as in Lemma 4.21, $\exists x' \between x\, \forall x'' \between x'\ \exists y''$: $x''Ry'' \between z$. Since $y'' \between z$ implies $y'' \notin A$ given $z \in \neg A$, it follows that $\exists x' \between x\, \forall x'' \between x'\ \exists y'' \in R(x'')$: $y'' \notin A$ and hence $x'' \notin \Box A$. ☐

The underlying compatibility frame of a modal compatibility frame gives rise to an ortholattice just as in Proposition 4.8, and the R relation gives us the modal operation of a modal ortholattice.

Proposition 4.23. For any modal compatibility frame $\mathcal{F} = \langle S, \between, i\rangle$, the $\between$-regular sets ordered by inclusion form a complete lattice, which becomes an ortholattice with $\neg$ defined as in (I) and then a modal ortholattice with $\Box$ as defined in Definition (III). We denote this modal ortholattice by $O(\mathcal{F})$.

Proof. Given Proposition 4.8, we need only check the new modal part, which is completely standard:

$$\begin{aligned}\Box(A \cap B) &= \{x \in S \mid R(x) \subseteq A \cap B\}\\ &= \{x \in S \mid R(x) \subseteq A\} \cap \{x \in S \mid R(x) \subseteq B\}\\ &= \Box A \cap \Box B,\end{aligned}$$

and $\Box 1 = \Box S = \{x \in S \mid R(x) \subseteq S\} = S = 1$. ☐

Now to interpret $R(x)$ as an *information state*, we impose two additional conditions on R. The first is the familiar reflexivity condition on epistemic accessibility.

Definition 4.24. A T *compatibility frame* is a modal compatibility frame satisfying:

- Reflexivity: for all $x \in S$, xRx.

As usual, it follows from this condition that $\Box A$ entails A.[21]

Proposition 4.25. For any T compatibility frame $\mathcal{F}$ and $\between$-regular set A, we have $\Box A \subseteq A$. Hence $O(\mathcal{F})$ is a T modal ortholattice.

Proof. If $x \in \Box A$, then $R(x) \subseteq A$, which with xRx implies $x \in A$. ☐

The second condition essentially says that *there is a possibility where everything settled true by x is known.*

Definition 4.26. An *epistemic compatibility frame* is a T compatibility frame also satisfying:

- Knowability: for all $x \in S$, there is a $y \in S$ such that for all $z \in R(y)$, $z \sqsubseteq x$.

Recall from Lemma 4.4 that $z \sqsubseteq x$ implies that every proposition true at x is also true at z.

Knowability can naturally be seen as the semantic constraint corresponding to the proof theoretic principle, noted in Lemma 3.25, that for any $\varphi \in \mathcal{EL}$, if $\Box\varphi \vdash_{\mathsf{EO}} \bot$, then $\varphi \vdash_{\mathsf{EO}} \bot$, or contrapositively, if $\varphi \nvdash_{\mathsf{EO}} \bot$, then $\Box\varphi \nvdash_{\mathsf{EO}} \bot$. In the presence of our background logic, that principle is equivalent to Wittgenstein's Law. Correspondingly, Knowability ensures that $\neg A$ and $\Diamond A$ are inconsistent in any epistemic compatibility frame.

Proposition 4.27. For any epistemic compatibility frame $\mathcal{F}$ and $\between$-regular set A, we have $\neg A \cap \Diamond A = \varnothing$. Hence $O(\mathcal{F})$ is an epistemic ortholattice.

Proof. Suppose $x \in \neg A$. By Knowability, there is a $y \in S$ such that for all $z \in R(y)$, we have $z \sqsubseteq x$ and hence $z \in \neg A$ by Lemma 4.4, so $y \in \Box\neg A$. Since $y \in R(y)$ by Reflexivity, it follows that $y \sqsubseteq x$ and hence $y \between x$, which with $y \in \Box\neg A$ implies $x \notin \neg\Box\neg A$, so $x \notin \Diamond A$. ☐

In essence, the explanation of the badness of Wittgenstein sentences according to our possibility semantics is this: since for any possibility x, *there is a possibility where everything settled true by x is known*, then x cannot settle $\neg A \wedge \Diamond A$ as true, for that would imply the existence of a possibility y settling as true the contradictory proposition $\Box(\neg A \cap \Diamond A) \subseteq \Box\neg A \cap \Box\Diamond A \subseteq \Box\neg A \cap \Diamond A = \varnothing$.

Remark 4.28. Some (e.g., Lassiter 2016) reject the T axiom for $\Box$. Without T, we can validate Wittgenstein's Law and its generalization in Fact 3.28 by adopting a slightly stronger constraint than Knowability: for all $x \in S$ and $n \in \mathbb{N}$, there is a $y \between x$ such that for all $z \in R^n(y)$, we have $z \sqsubseteq x$, where $R^0(y) = \{y\}$ and $R^{n+1}(y) = R[R^n(y)]$. Clearly this suffices for the proof in Proposition 4.27 that $\neg A \cap \Diamond A = \varnothing$ and its generalization to $\neg A \cap \Diamond^n A = \varnothing$, so those who reject T can still use our approach to account for the contradictoriness of (generalized) Wittgenstein sentences.

Remark 4.29. When we say 'there is a possibility where everything settled true by x is known', we are using 'known' in a loose way. If we identified epistemic modals with the knowledge of some particular agent (or group) at a time—the speaker, say—then this condition would be somewhat implausible in general, since we can plausibly imagine possibilities where it is settled that not everything settled is known. Instead, we are thinking about 'what is known' as corresponding to something like salient information (possibly) augmented with further facts that may well go beyond what any individual or group knows. The idea that epistemic modals can track information that goes beyond what is actually known by anyone goes back to Hacking (1967). Hacking considers

[21] In fact, the following weaker condition suffices: for all $x \in S$, there is a $y \in R(x)$ with $x \sqsubseteq y$.

a salvage crew searching for a ship that sank a long time ago. The mate of the salvage ship works from an old log, makes some mistakes in his calculations, and concludes that the wreck may be in a certain bay. It is possible, he says, that the hulk is in these waters. No one knows anything to the contrary. But in fact, as it turns out later, it simply was not possible for the vessel to be in that bay; more careful examination of the log shows that the boat must have gone down at least thirty miles further south. The mate said something false when he said, 'It is possible that we shall find the treasure here,' but the falsehood did not arise from what anyone actually knew at the time. (Hacking, 1967, p. 148)

(We can, if we want, elaborate stories like this so that the falsehood does not arise from what anyone actually knows at *any* time.) This idea is elaborated in different ways in Dorr and Hawthorne 2013, Mandelkern 2019, and Kratzer 2020, and we think it is fundamentally correct. Of course, this idea remains vague and context-sensitive; nonetheless, we think that our Knowability constraint is a natural way of making precise the idea that epistemic modalities involve some mix of purely epistemic and factual considerations.

To summarize, we impose the following three conditions:

- R-regularity: if $xRy' \between y$, then $\exists x' \between x$: $x' \in \Diamond{\downarrow}y$. (if x can epistemically access a possibility compatible with y, then x is compatible with a possibility according to which y *might obtain*);
- Reflexivity: for all $x \in S$, xRx
  (everything known is true);
- Knowability: for all $x \in S$, there is a $y \in S$ such that for all $z \in R(y)$, we have $z \sqsubseteq x$
  (there is a possibility where everything settled true by x is known).

Example 4.30. Figure 12 shows an epistemic compatibility frame $\mathcal{F}$ based on the compatibility frame of Example 4.11. It is a good exercise to check that the three conditions on the R relation are satisfied. The epistemic ortholattice $O(\mathcal{F})$ is isomorphic to the epistemic ortholattice in Figure 3, as we will see when we discuss a model based on the frame $\mathcal{F}$ in Example 4.33.

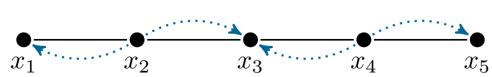

Figure 12: An epistemic compatibility frame with dotted arrows indicating the accessibility relation R (with reflexive loops omitted)

Turning to formal semantics for $\mathcal{EL}$ (recall Definition 3.21), Proposition 4.23 leads to the following.

Definition 4.31. Given a modal compatibility model $\mathcal{M} = \langle S, \between, i, V\rangle$, $x \in S$, and $\varphi \in \mathcal{L}$, we define $\mathcal{M}, x \Vdash \varphi$ with the same clauses as in Definition 4.15 plus

- $\mathcal{M}, x \Vdash \Box\varphi$ iff for all $y \in R(x)$, we have $\mathcal{M}, y \Vdash \varphi$.

As before, we define $[\![\varphi]\!]^{\mathcal{M}} = \{x \in S \mid \mathcal{M}, x \Vdash \varphi\}$.

Then given our definition of $\Diamond$ as $\neg\Box\neg$, as in (IV) we have:

- $\mathcal{M}, x \Vdash \Diamond\varphi$ iff $\forall x' \between x\ \exists y' \in R(x')\ \exists y'' \between y'$: $\mathcal{M}, y'' \Vdash \varphi$.

Once again an easy induction shows the following.

Lemma 4.32. For any modal compatibility model $\mathcal{M}$ and $\varphi \in \mathcal{EL}$, $[\![\varphi]\!]^{\mathcal{M}}$ is a $\between$-regular set.

Example 4.33. Figure 13 shows an epistemic compatibility model based on the epistemic compatibility frame in Figure 12 with $V(p) = \{x_1, x_2\}$, so $V(p)$ is a $\between$-regular set and $[\![\neg p]\!]^{\mathcal{M}} = \{x_4, x_5\}$. We call this model the Epistemic Scale for p, since as we move from x_1 to x_5 we move as if on a scale from $\Box p$ at x_1 to p at x_2 to $\Diamond p \wedge \Diamond \neg p$ at x_3 to $\neg p$ at x_4 to $\Box \neg p$ at x_5. Observe the following:

- $[\![\Box p]\!]^{\mathcal{M}} = \{x_1\}$;
- $[\![\neg \Box p]\!]^{\mathcal{M}} = [\![\Diamond \neg p]\!]^{\mathcal{M}} = \{x_3, x_4, x_5\}$;
- $[\![\Box \neg p]\!]^{\mathcal{M}} = \{x_5\}$;
- $[\![\neg \Box \neg p]\!]^{\mathcal{M}} = [\![\Diamond p]\!]^{\mathcal{M}} = \{x_1, x_2, x_3\}$;
- $[\![\Diamond p \wedge \Diamond \neg p]\!]^{\mathcal{M}} = \{x_3\}$;
- $[\![\Box p \vee \Box \neg p]\!]^{\mathcal{M}} = \{x_1, x_5\}$.

These calculations match the fact that the associated epistemic ortholattice is isomorphic to that in Figure 3.

Let us also see how our possibility semantics explains the failure of the distributive law of classical logic. Consider x_3, which is a *partial* possibility: it does not settle p as true and it does not settle $\neg p$ as true, as it is compatible with both (since it is compatible with x_2 and x_4), but like any possibility, it settles $p \vee \neg p$ as true. Now since the information available in x_3 is x_3 itself, x_3 settles that p might be true and that $\neg p$ might be true: $\Diamond p \wedge \Diamond \neg p$. Given that $(p \vee \neg p) \wedge (\Diamond p \wedge \Diamond \neg p)$ is true at x_3, the distributive law would require that $(p \wedge \Diamond \neg p) \vee (\neg p \wedge \Diamond p)$ be true at x_3. But we have already seen in our discussion after Proposition 4.27 why no possibility can settle either of those disjuncts as true. Then since no possibility settles either disjunct as true, x_3 does not settle the disjunction as true, which shows that the distributive law is invalid.

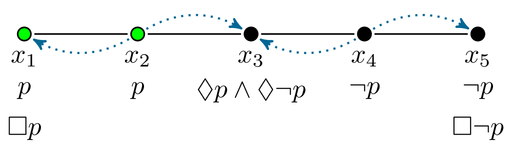

Figure 13: An epistemic compatibility model, dubbed the Epistemic Scale for p, based on the frame of Figure 12, with $V(p)$ highlighted in green.

Example 4.34. Let us consider an example with two dimensions of epistemic variation in contrast to the one dimension of the Epistemic Scale. Figure 14 shows what we call an Epistemic Grid for p and q. It is obtained from the Epistemic Scale by a product construction: each possibility in the Epistemic Grid is a pair of possibilities (x, y) with x from the Epistemic Scale for p and y from the Epistemic Scale for q; (x, y) is compatible with (resp. accessible from) (x', y') iff x is compatible with (resp. accessible from) x' and y is compatible with (resp. accessible from) y';[22] and a propositional variable is true at (x, y) iff it is true at x or at y. We encourage the reader to calculate the truth values of some formulas at possibilities in this Epistemic Grid. For example, $\neg(p \wedge q)$ is true at the possibilities in the bottommost two rows and rightmost two columns; hence $\Box \neg (p \wedge q)$ is true in the bottommost row and rightmost column; and hence $\neg \Box \neg (p \wedge q)$, i.e., $\Diamond(p \wedge q)$, is true at the nine possibilities in the upper-left quadrant, which are also the possibilities where $\Diamond p \wedge \Diamond q$ is true. By contrast, along the centermost column, $\Diamond p \wedge \Diamond \neg p \wedge \neg \Diamond (p \wedge \neg p)$ is true.

While the epistemic ortholattice arising from an Epistemic Scale has 10 elements, shown on the right of Figure 10, the epistemic ortholattice arising from an Epistemic Grid has 1,942 elements.[23] This is another

[22] This construction of the accessibility relation guarantees that R-regularity, Reflexivity, and Knowability are still satisfied.

[23] Though that is small compared to the size of the powerset of the Epistemic Grid at $2^{25} = 33,554,432$ elements. See the second of the notebooks cited in § 1 for the calculation of the 1,942 figure.

reason why we want a possibility semantics in addition to an algebraic semantics: while a 1,942-element epistemic ortholattice is too large to grasp, we can represent it using a perspicuous 25-element relational epistemic compatibility frame. One can even visualize an Epistemic Cube for three variables and consider more general n-dimensional Epistemic Hypercubes (see the second notebook cited in § 1).

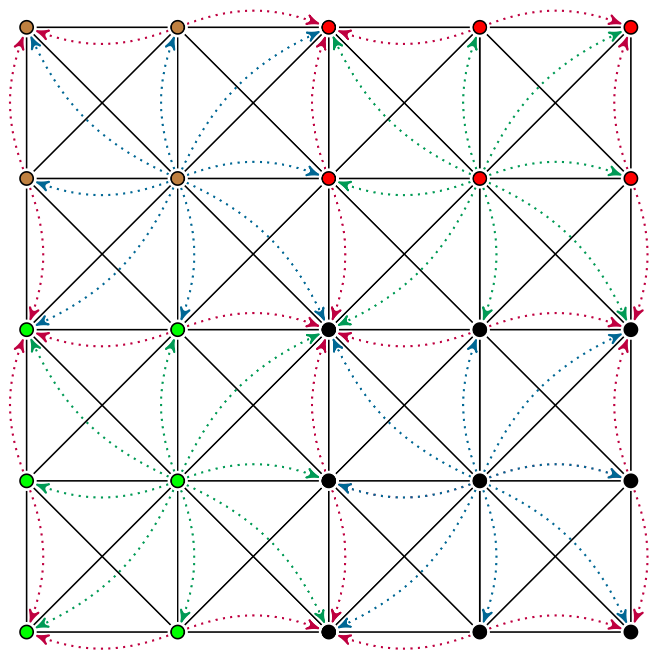

Figure 14: An Epistemic Grid for p and q. Green possibilities make p true; red possibilities make q true; and brown possibilities makes both p and q true. Different colors for different dotted arrows are used only to make the pattern more intelligible; all such arrows represent the accessibility relation. Reflexive dotted loops between each possibility and itself are assumed but omitted from the diagram.

Remark 4.35. While we interpret the 'must' modality $\Box$ using the R relation, there is another operation, which we will denote by $\Box_{\between}$, available on the ortholattice $O(\mathcal{F})$ of any compatibility frame $\mathcal{F}$: $\Box_{\between} A = \{x \in S \mid \forall x' \between x \; x' \in A\}$. Note that if A is $\between$-regular, so is $\Box_{\between} A$: for if $x \notin \Box_{\between} A$, then there is a $y \between x$ with $y \notin A$, which implies that for all $z \between y$, $z \notin \Box_{\between} A$.[24]

In fact, interpreting 'must' as $\Box_{\between}$ is just a special case of our possibility semantics: the proposal is just to set $R = \between$. However, while the $\Box_{\between}$ modality may have useful applications,[25] it is not appropriate for interpreting 'must'. For one thing, it validates the B principle (where $\Diamond_{\between} A = \neg\Box_{\between}\neg A$): for all $A \in O(\mathcal{F})$, $\Diamond_{\between}\Box_{\between} A \subseteq A$, which we do not want to require in general. Even in cases where B is acceptable, $\Box_{\between}$ may

[24]Note, by contrast, that the function $\blacklozenge_{\between}$ defined by $\blacklozenge_{\between} A = \{x \in S \mid \exists x' \between x : x' \in A\}$ is not an operation on the ortholattice $O(\mathcal{F})$, i.e., it is not guaranteed that if A is $\between$-regular, then so is $\blacklozenge_{\between} A$. For example, in the compatibility frame in Figure 7, $\{x_1, x_2, x_3\}$ is $\between$-regular, but $\blacklozenge_{\between}\{x_1, x_2, x_3\} = \{x_1, x_2, x_3, x_4\}$ is not. A natural response to this problem is to apply the closure operator from Footnote 18 to $\blacklozenge_{\between} A$, interpreting 'might A' as $c_{\between}\blacklozenge_{\between} A$. But $c_{\between}\blacklozenge_{\between}$ is just the dual of $\Box_{\between}$, i.e., $c_{\between}\blacklozenge_{\between} A = \neg\Box_{\between}\neg A$, so the reasons for rejecting $\Box_{\between}$ as 'must' in the remainder of Remark 4.35 are also reasons for rejecting $c_{\between}\blacklozenge_{\between}$ as 'might'.

[25]For example, compatibility frames can of course also be regarded as classical possible world frames with $\between$ as an accessibility relation for the modality $\Box_{\between}$ on the powerset of the set of possibilities. This leads to the translation of orthologic into the classical modal logic KTB (Goldblatt 1974, cf. Dishkant 1977), similar in spirit to the translation of intuitionistic logic into classical S4 (Gödel 1933, McKinsey and Tarski 1948). Likewise, epistemic compatibility frames can be regarded as classical *bimodal* possible world frames, just as possibility frames for classical modal logic are regarded as bimodal possible world frames in van Benthem et al. 2017. This leads to a translation of our epistemic orthologic into a classical bimodal logic. We omit the details, which can be worked out on the model of van Benthem et al. 2017.

not give the desired results for 'must'. For example, in the frame in Figure 13, where $P = \{x_1, x_2\}$, we have $\Box P = \Box_{\between} P = \{x_1\}$, $\Box\neg P = \Box_{\between}\neg P = \{x_5\}$, and $\neg\Box P \wedge \neg\Box\neg P = \neg\Box_{\between} P \wedge \neg\Box_{\between}\neg P = \{x_3\}$; however, while $\Box\Box P = \{x_1\}$, $\Box\neg\Box P = \{x_5\}$, and $\Box(\neg\Box P \wedge \neg\Box\neg P) = \{x_3\}$, we have $\Box_{\between}\Box_{\between} P = \Box_{\between}\neg\Box_{\between} P = \Box_{\between}(\neg\Box_{\between} P \wedge \neg\Box_{\between}\neg P) = \varnothing$. The general problem is that iterating $\Box_{\between}$ keeps shrinking the relevant set until one reaches a set of possibilities that are not compatible with anything outside it. Most fatal for the idea of interpreting 'must' as $\Box_{\between}$ is that this approach fails to make Wittgenstein sentences contradictions. Consider a compatibility frame with three possibilities such that $x \between y \between z$, but not $x \between z$. Set $V(p) = \{x\}$, which is a $\between$-regular set. Then $p \wedge \neg\Box_{\between} p$ is true at x. The moral is that we cannot reduce *epistemic* modality entirely to the *alethic* relation of compatibility.

As before, we define semantic consequence standardly in terms of truth preservation.

Definition 4.36. Given a class $\mathbf{C}$ of modal compatibility frames, define the semantic consequence relation $\vDash_{\mathbf{C}}$, a binary relation on $\mathcal{EL}$, as follows: $\varphi \vDash_{\mathbf{C}} \psi$ if for every $\mathcal{F} \in \mathbf{C}$, model $\mathcal{M}$ based on $\mathcal{F}$, and possibility x in $\mathcal{M}$, if $\mathcal{M}, x \Vdash \varphi$, then $\mathcal{M}, x \Vdash \psi$.

Our first main result is the completeness of the epistemic orthologic EO in Definition 3.24 with respect to our epistemic possibility semantics.

Theorem 4.37. The logic EO is sound and complete with respect to the class **ECF** of all epistemic compatibility frames according to the consequence relation of Definition 4.36: for all $\varphi, \psi \in \mathcal{EL}$, we have $\varphi \vdash_{\mathsf{EO}} \psi$ if and only if $\varphi \vDash_{\mathbf{ECF}} \psi$.

Proof. Soundness is a straightforward check of the principles of the logic. For completeness, as in the proofs of Theorem 3.13 and 3.26, we consider the Lindenbaum-Tarski algebra L of EO, which is an epistemic ortholattice. From L we define a modal compatibility frame $\mathcal{F} = \langle S, \between, R\rangle$ as follows:

- S is the set of all proper filters of L;[26]
- for $F, G \in S$, $F \between G$ if there is no $a \in F$ with $\neg a \in G$;
- for $F, G \in S$, FRG iff for all $\Box a \in F$, we have $a \in G$.

That $\between$ is reflexive follows from the fact that each $F \in S$ is a *proper* filter, and in L, $a \wedge \neg a = 0$. That $\between$ is symmetric follows from the fact that if $a \in F$ and $\neg a \in G$, then $\neg a \in G$ and $\neg\neg a \in F$, since $a = \neg\neg a$ in L, so there is a b such that $b \in G$ and $\neg b \in F$. That R is reflexive follows from L being a T modal ortholattice.

Below we will use the fact that if F is a proper filter, then so is $i(F) = \{a \in B \mid \Box a \in F\}$. For if not, then there are $\Box a_1, \dots, \Box a_n \in F$ such that $a_1 \wedge \dots \wedge a_n = 0$, which implies $\Box(a_1 \wedge \dots \wedge a_n) = \Box 0$ and hence $\Box a_1 \wedge \dots \wedge \Box a_n = \Box 0$. Since L is a T modal ortholattice, $\Box 0 = 0$, so we obtain $\Box a_1 \wedge \dots \wedge \Box a_n = 0$, which contradicts the assumption that F is proper, since $\Box a_1, \dots, \Box a_n \in F$ implies $\Box a_1 \wedge \dots \wedge \Box a_n \in F$.

Now let us prove R-regularity in its form in Lemma 4.21:

- if $FRG' \between G$, then $\exists F' \between F\ \forall F'' \between F'\ \exists G''$: $F''RG'' \between G$.

Let F' be the filter generated by $\{\Diamond a \mid a \in G\}$. We claim that F' is proper and that $F' \between F$. If not, then for some $b \in F'$, $\neg b \in F$. Since $b \in F'$, there are $a_1, \dots, a_n \in G$ such that $\Diamond a_1 \wedge \dots \wedge \Diamond a_n \leq b$. It follows that $\Diamond(a_1 \wedge \dots \wedge a_n) \leq b$, so $\neg b \leq \Box\neg(a_1 \wedge \dots \wedge a_n)$. Hence $\Box\neg(a_1 \wedge \dots \wedge a_n) \in F$, which with FRG' implies

[26] Recall that a *filter* in a lattice is a nonempty set F of elements such that for any $a, b \in L$, together $a \in F$ and $a \leq b$ imply $b \in F$, and $a, b \in F$ implies $a \wedge b \in F$; a filter is *proper* if it does not contain every element of the lattice.

$\neg(a_1 \wedge \cdots \wedge a_n) \in G'$. But $a_1, \ldots, a_n \in G$ implies $a_1 \wedge \cdots \wedge a_n \in G$, so $\neg(a_1 \wedge \cdots \wedge a_n) \in G'$ contradicts $G' \between G$. Thus, we have $F' \between F$. Now consider any $F'' \between F'$. We claim that $i(F'') \between G$. If not, then there is some $a \in G$ such that $\neg a \in i(F'')$. But $a \in G$ implies $\Diamond a \in F'$, and $\neg a \in i(F'')$ implies $\Box\neg a \in F''$, and together $\Diamond a \in F'$ and $\Box\neg a \in F''$ contradict $F'' \between F'$. Thus, $i(F'') \between G$, so setting $G'' = i(F'')$ completes the proof of R-regularity.

Finally, we show that the frame satisfies Knowability: for all $F \in S$, there is some $G \in S$ such that for all $H \in R(G)$, we have $H \sqsubseteq F$. Given F, let G be the filter generated by $\{\Box a \mid a \in F\}$. We claim that G is proper. If not, then there are $a_1, \ldots, a_n \in F$ such that $\Box a_1 \wedge \cdots \wedge \Box a_n = 0$ and hence $\Box(a_1 \wedge \cdots \wedge a_n) = 0$. Hence $\Diamond\neg(a_1 \wedge \cdots \wedge a_n) = 1$, so $\Diamond\neg(a_1 \wedge \cdots \wedge a_n) \in F$. From $a_1, \ldots, a_n \in F$ we also have $a_1 \wedge \cdots \wedge a_n \in F$. But since L is an epistemic ortholattice, from $\Diamond\neg(a_1 \wedge \cdots \wedge a_n) \in F$ and $a_1 \wedge \cdots \wedge a_n \in F$, we have $0 \in F$, contradicting our assumption that F is proper. Hence G is proper. To show $H \sqsubseteq F$, suppose $I \between H$ but not $I \between F$, so there is some $b \in I$ such that $\neg b \in F$. Since $\neg b \in F$, $\Box\neg b \in G$, which with GRH implies $\neg b \in H$, which contradicts $I \between H$.

Thus, $\mathcal{F} = \langle S, \between, i\rangle \in \mathbf{ECF}$. It is well known that the map sending $a \in L$ to $\widehat{a} = \{F \in S \mid a \in F\}$ is an embedding of the ortholattice reduct of L into the ortholattice of $\between$-regular subsets of $\langle S, \between\rangle$ (see, e.g., Goldblatt 1975, Proposition 1).[27] It only remains to observe that this embedding also respects $\Box$:

$$\widehat{\Box a} = \{F \in S \mid \Box a \in F\} = \{F \in S \mid a \in i(F)\} = \{F \in S \mid R(F) \subseteq \widehat{a}\} = \Box\widehat{a}.$$

Thus, the map $a \mapsto \widehat{a}$ is a modal ortholattice embedding of L into $O(\mathcal{F})$. Let $\mathcal{M}$ be the model based on $\mathcal{F}$ with the valuation V defined by $V(p) = \widehat{[p]}$. Then since $a \mapsto \widehat{a}$ is a modal ortholattice embedding, we have $[\![\varphi]\!]^{\mathcal{M}} = \widehat{[\varphi]}$ for all $\varphi \in \mathcal{EL}$. Now if $\varphi \nvdash_{\mathsf{EO}} \psi$, then in L we have $[\varphi] \nleq [\psi]$. Then since $a \mapsto \widehat{a}$ is a modal ortholattice embedding, it follows that $\widehat{[\varphi]} \nsubseteq \widehat{[\psi]}$ and hence $[\![\varphi]\!]^{\mathcal{M}} \nsubseteq [\![\psi]\!]^{\mathcal{M}}$. Therefore, $\varphi \nvDash_{\mathbf{ECF}} \psi$. □

Remark 4.38. The proof of Theorem 4.37 shows that instead of giving semantics for $\Box$ using an accessibility relation $R \subseteq S \times S$, we could give a semantics using an accessibility *function* $i : S \to S$ such that

$$\Box A = \{x \in S \mid i(x) \in A\}.$$

This semantics simplifies some abstract calculations but has the disadvantage of requiring us to add more possibilities to frames such as the frame in Figure 12 in order to do the work of a set of epistemically accessible possibilities using a single coarse-grained possibility (in the case of the frame in Figure 12, the functional approach requires us to add two possibilities to the frame).

4.3 Distinguishing non-epistemic propositions

Just as we distinguished non-epistemic propositions from arbitrary propositions in the context of algebraic semantics in § 3.2.4, we can now do the same in the context of possibility semantics.[28]

Proposition 4.39. Given a compatibility frame $\mathcal{F} = \langle S, \between, R\rangle$, let $\mathbb{B}$ be a nonempty set of $\between$-regular subsets of S closed under $\cap$ and the operation $\neg$ from (I). Then the following are equivalent:

[27] Goldblatt (1975) works with the complement of the compatibility relation, called the orthogonality relation, but the proof is easily rephrased in terms of compatibility.

[28] Equipping a frame with a distinguished collection of propositions is similar in spirit to introducing *general frames* in modal logic (Blackburn et al. 2001, § 5.5). Unlike with general frames, we do not require that the distinguished collection be closed under $\Box$. However, we will have a different requirement involving $\Box$ in Definition 4.40.

1. $\mathbb{B}$ is a Boolean algebra under $\cap$ and $\neg$;
2. for all $A, B \in \mathbb{B}$, if there are $x \in A$ and $y \in B$ with $x \between y$, then $A \cap B \neq \varnothing$.

Proof. By Proposition 3.7, $\mathbb{B}$ being Boolean is equivalent to the condition that $\neg$ restricted to $\mathbb{B}$ is pseudo-complementation, i.e., that for all $A, B \in \mathbb{B}$, if $A \cap B = \varnothing$, then $A \subseteq \neg B$; and $A \subseteq \neg B$ is equivalent to there being no $x \in A$ and $y \in B$ with $x \between y$. Contraposing, we obtain the condition in part 2. ☐

Thus, the key condition on $\mathbb{B}$ is that if there are compatible possibilities making A and B true, respectively, then there is a single possibility making both A and B true (cf. the even stronger classical idea that compatibility implies compossibility in Remark 4.9).

Definition 4.40. A *grounded modal compatibility frame* is a tuple $\mathcal{F} = \langle S, \between, R, \mathbb{B} \rangle$ where $\langle S, \between, R \rangle$ is a modal compatibility frame and $\mathbb{B}$ is a nonempty collection of $\between$-regular sets closed under $\cap$ and the operation $\neg$ from (I). Given a grounded frame $\mathcal{F} = \langle S, \between, R, \mathbb{B} \rangle$,

- let $\mathbb{B}_0 = \mathbb{B}$, and
- let $\mathbb{B}_{n+1}$ be the closure of $\{\Box B \mid B \in \mathbb{B}_n\}$ under $\cap$ and $\neg$,

where $\Box$ is the operation defined from R in (III). We say that $\mathcal{F}$ is *stratified* if for all $n \in \mathbb{N}$ and $A, B \in \mathbb{B}_n$, if there are $x \in A$ and $y \in B$ with $x \between y$, then $A \cap B \neq \varnothing$.

A *stratified epistemic compatibility frame* is a stratified modal compatibility frame in which $\langle S, \between, R \rangle$ is an epistemic compatibility frame.

Then the following is immediate from Proposition 4.39.

Corollary 4.41. In a stratified modal compatibility frame, each $\mathbb{B}_n$ is a Boolean algebra.

Example 4.42. It is easy to check by hand that the epistemic compatibility frame in Figure 12 equipped with $\mathbb{B} = \{\varnothing, \{x_1, x_2\}, \{x_4, x_5\}, S\}$ is stratified. In this case $\mathbb{B}_1$ has eight propositions, and $\mathbb{B}_2 = \mathbb{B}_1$, so $\mathbb{B}_n = \mathbb{B}_1$ for all $n \geq 1$; this matches what we observed in Example 3.33 for the corresponding epistemic ortho-Boolean lattice.

Recall from § 3.2.4 the distinguished set Bool of non-epistemic propositional variables. These are now interpreted as propositions in the distinguished family $\mathbb{B}$.

Definition 4.43. A *grounded modal compatibility model* is a pair $\mathcal{M} = \langle \mathcal{F}, V \rangle$ where $\mathcal{F} = \langle S, \between, R, \mathbb{B} \rangle$ is a grounded modal compatibility frame and V assigns to each $p \in \mathsf{Prop}$ a $\between$-regular $V(p) \subseteq S$ and to each $\mathsf{p} \in \mathsf{Bool}$ a $V(\mathsf{p}) \in \mathbb{B}$. We say that $\mathcal{M}$ is *based on* $\mathcal{F}$.

Next comes the definition of consequence, which follows exactly the same pattern as before.

Definition 4.44. Given a class $\mathbf{C}$ of grounded modal compatibility frames, define the semantic consequence relation $\vDash_{\mathbf{C}}$, a binary relation on $\mathcal{EL}^+$, as follows: $\varphi \vDash_{\mathbf{C}} \psi$ if for every $\mathcal{F} \in \mathbf{C}$, model $\mathcal{M}$ based on $\mathcal{F}$, and possibility x in $\mathcal{M}$, if $\mathcal{M}, x \Vdash \varphi$, then $\mathcal{M}, x \Vdash \psi$.

Finally, we can prove that our epistemic orthologic EO^+ in Figure 5 is sound and complete with respect to possibility semantics using stratified epistemic compatibility frames.

Theorem 4.45. The logic EO^+ is sound and complete with respect to the class $\mathbf{ECF}^+$ of all stratified epistemic compatibility frames according to the consequence relation of Definition 4.44: for all $\varphi, \psi \in \mathcal{EL}^+$, we have $\varphi \vdash_{\mathsf{EO}^+} \psi$ if and only if $\varphi \vDash_{\mathbf{ECF}^+} \psi$.

Proof. Simply add to the proof of Theorem 4.37 that $\mathbb{B} = \{\widehat{[\beta]} \mid \beta \text{ Boolean}\}$. As in the proof of Theorem 3.40, each B_n is a Boolean subalgebra of the Lindenbaum-Tarski algebra, which implies that $\mathbb{B}_n$ satisfies the key condition in Definition 4.40: in contrapositive form, if there is no proper filter $F \in \widehat{[\alpha]} \cap \widehat{[\beta]}$, i.e., no proper filter F with $[\alpha] \wedge [\beta] \in F$, so $[\alpha] \wedge [\beta] = 0$ in the Lindenbaum-Tarski algebra, then since $[\alpha]$ and $[\beta]$ belong to a Boolean subalgebra, we have $[\alpha] \leq \neg[\beta]$, which implies that there are no proper filters $F \in \widehat{[\alpha]}$ and $G \in \widehat{[\beta]}$ with $F \between G$. □

Thus, we have now reached our "mission accomplished" moment: an E-logic, satisfying all the desiderata of § 2, drops out of a concrete possibility semantics generalizing possible world semantics.

5 Constructing possibilities from worlds

We have now given an algebraic semantics and a possibility semantics that both correspond to our target logic EO^+. From one perspective, our job is now done twice over: for those who might worry that algebraic semantics is suspiciously close to the corresponding syntactic principles, we have given a concrete implementation in terms of truth at possibilities. However, some might find it hard to theorize directly in terms of possibilities. Possible worlds are generally more familiar to semanticists and logicians, and they have been much more thoroughly studied by philosophers.

In order to counteract any squeamishness about possibilities, we will show in this section that we can construct a model for our possibility semantics from any given possible worlds model. This construction should put to rest any worries that readers might have about working with a framework that rests on a non-standard semantic foundation. It shows that our approach is ontologically innocent: if you are comfortable with possible worlds but not possibilities, you can build the whole system on that foundation. (Of course, there may be ontological reasons to prefer possibilities, but we will not make that argument here.) Moreover, by giving a concrete way for those familiar with possible worlds semantics to build toy models in our semantic framework, this construction will enable semanticists comfortable with possible worlds to easily build toy models for our possibility semantics.

In more detail, given a possible worlds model, we will show how to transform that model into one of our epistemic compatibility models. As we will see, the resulting models validate a strict strengthening of EO^+; in particular, they commit us to the $\mathsf{S5}$ axioms in addition to EO^+, so not every epistemic compatibility model can be represented as coming from a possible worlds model. The goal of this construction is to show just one way to obtain a possibility semantics from a non-modal starting point.

The basic idea behind the construction, starting with a set W of worlds, is to

construct possibilities as *pairs* (A, I) *of sets of worlds* where $\varnothing \neq A \subseteq I \subseteq W$.

In fact, we show in § 5.1 that our construction applies starting with an arbitrary Boolean algebra B, but in the finite case we can assume without loss of generality that B is the powerset of a set of worlds.

Moreover, this construction gives us a general scaffolding for moving from structures defined on possible worlds to structures defined on possibilities. In particular, we will show in § 5.2 how a probability function on sets of worlds can be extended to a probability function on our epistemic extension of the powerset algebra; and in § 5.3 we show how a conditional selection function defined on world-proposition pairs can be extended to a conditional selection function defined on possibility-proposition pairs. In both cases, we will briefly point to some virtues of the resulting construction, though we defer extensive discussion to future work. Our present goal is to show how this lifting technique can nimbly take us from the familiar setting of possible worlds semantics into a possibility semantics framework more suitable for theorizing about epistemic modality—and thus how our possibility semantics can capture both the standard phenomena covered by possible world semantics and the peculiar behavior of epistemic modals.

5.1 Epistemic extension

We begin with the basic construction that takes us from a possible worlds model—or, more generally, any Boolean algebra of propositions—to a possibility model. Given a Boolean algebra B of non-modal propositions, we wish to "add modal propositions" to B, resulting in an epistemic ortholattice that we will call the *epistemic extension* of B. We will define the epistemic extension of B as the epistemic ortholattice coming from an epistemic compatibility frame constructed from B.

Definition 5.1. Let B be a Boolean algebra. The *epistemic frame of* B is the tuple $B^{\mathsf{e}} = (S, \between, R)$ defined as follows:

1. $S = \{(a, i) \mid a, i \in B, 0 \neq a \leq i\}$;
2. $(a, i) \between (a', i')$ iff $a \wedge a' \neq 0$ and $a \leq i'$ and $a' \leq i$;
3. $(a, i)R(a', i')$ iff $a \leq a'$ and $i' \leq i$.

Given a valuation $\theta : \mathsf{Bool} \to B$, we define θ^{e} by $\theta^{\mathsf{e}}(\mathtt{p}) = \{(a, i) \mid a \leq \theta(\mathtt{p})\}$. The *epistemic model of* (B, θ) is the pair $(B^{\mathsf{e}}, \theta^{\mathsf{e}})$.

Thus, the possibilities in the epistemic frame of B are pairs of propositions. As noted above, if B is the powerset of a set W of worlds, the possibilities are pairs (A, I) where $\varnothing \neq A \subseteq I \subseteq W$. Pairs of the form $(\{w\}, I)$ are familiar from other semantics for epistemic modals (e.g., Yalcin 2007, MacFarlane 2011), where $w \in W$ is a world and $I \subseteq W$ is an information state. But note that our A need not be a singleton set, so we have in effect "partialized" or "possibilized" the worldly component of the familiar picture. Moreover, we shall see below (Lemmas 5.4 and 5.8) that while what *must* be the case is determined by the second component I, as in the familiar picture, what *might* be the case is determined by the same component A that determines what non-modal propositions *are* the case (at least when 'must' and 'might' apply to Boolean propositions), which gets to the heart of why our semantics validates Wittgenstein's Law.

In particular, any Boolean β *entailed* by the first component of a possibility (i.e., $A \subseteq \tilde{\theta}(\beta)$, when B is a powerset) *is* true, any Boolean β *consistent* with the first component (i.e., $A \cap \tilde{\theta}(\beta) \neq \varnothing$) *might* be true, and any Boolean β *entailed* by the second component (i.e., $I \subseteq \tilde{\theta}(\beta)$) *must* be true. This helps explains the definitions of $\between$ and R. The definition of $\between$ says that two possibilities are compatible if (i) how things are and might be according to each possibility is consistent with how things are and might be according to the other ($A \cap A' \neq \varnothing$) and (ii) how things are and might be according to each possibility is a narrowing down of how things must be according to the other ($A \subseteq I'$ and $A' \subseteq I$). While (i) ensures that if one possibility settles that β is true, then a compatible possibility does not settle that $\neg\beta$ is true, (ii) ensures that if one possibility settles that β *might* be true (so there is a β-world in A), then a compatible possibility does not settle that $\neg\beta$ *must* be true (since the β-world in A belongs to I'). According to the definition of R, one possibility has epistemic access to another iff both components of the second possibility lie in the interval $\{X \subseteq W \mid A \subseteq X \subseteq I\}$ between the two components of the first. This ensures that whatever β might be true according to the first possibility also might be true according to the second (since $A \subseteq A'$), and whatever β must be true according to the first possibility also must be true according to the second (since $I' \subseteq I$).

We can also give a simple characterization of refinement (recall Lemma 4.4) in the epistemic frame of B.

Lemma 5.2. Given a Boolean algebra B and possibilities (a, i) and (a', i') in B^{e}, we have $(a, i) \sqsubseteq (a', i')$ iff $a = a'$ and $i \leq i'$.

Proof. For the right-to-left direction, suppose $a = a'$, $i \leq i'$, and $(a'', i'') \between (a, i)$. Then since $a = a'$ and $i \leq i'$, we have $(a'', i'') \between (a', i')$. Hence $(a, i) \sqsubseteq (a', i')$. From left to right, if $a \not\leq a'$, then $(a \wedge \neg a', a) \between (a, i)$ but not $(a \wedge \neg a', a) \between (a', i')$, so $(a, i) \not\sqsubseteq (a', i')$. If $a' \not\leq a$, then $(a, a) \between (a, i)$ but not $(a, a) \between (a', i')$, so $(a, i) \not\sqsubseteq (a', i')$. Finally, if $i \not\leq i'$, then $(i, i) \between (a, i)$ but not $(i, i) \between (a', i')$, so $(a, i) \not\sqsubseteq (a', i')$. ☐

Below we will prove that B^{e} is indeed an epistemic compatibility frame (Definition 4.26), but first let us consider some concrete examples. Given a set W of worlds, we consider the epistemic frame of $\wp(W)$.

Example 5.3. Suppose we are about to flip a fair coin, so $W = \{0, 1\}$. Figure 15 shows the epistemic frame of $\wp(W)$. Note that it is isomorphic to the frame underlying the Epistemic Scale in Figure 13. As before, the solid lines represent compatibility (with reflexive loops omitted), and the dotted lines represent epistemic accessibility (with reflexive loops omitted). For example, according to Definition 5.1.2, the possibility $x_1 = (\{0\}, \{0\})$ is compatible with $x_2 = (\{0\}, \{0, 1\})$ because their first-coordinates have a non-empty intersection, and the first coordinate of each entails the second coordinate of the other. However, x_1 is not compatible with $x_3 = (\{0,\ 1\}, \{0, 1\})$, since the first coordinate of x_3 does not entail the second coordinate of x_1. As for accessibility, according to Definition 5.1.3, x_2 can access x_3 because both coordinates of x_3 lie in the interval between the coordinates of x_2. However, x_3 cannot access x_2 because the first coordinate of x_2 does not lie in the interval between the coordinates of x_3.

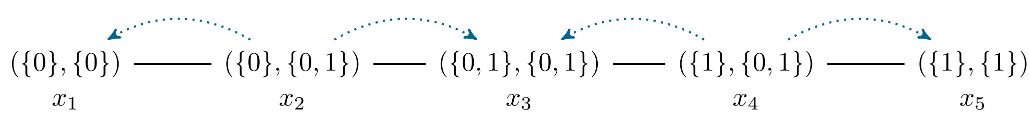

Figure 15: Epistemic frame constructed from two worlds.

An especially appealing aspect of this construction is how simple it is to check the truth values of formulas of modal depth ≤ 1 at a possibility (A, I), as shown by the following lemma. Given a possible worlds model $\mathfrak{M} = (W, V)$, where W is a nonempty set and $V : \mathsf{Bool} \to \wp(W)$, and a Boolean formula φ (i.e., one free of modals, as in Definition 3.29), define $\mathfrak{M}, w \vDash \varphi$ as usual: $\mathfrak{M}, w \vDash \mathtt{p}$ iff $w \in V(\mathtt{p})$; $\mathfrak{M}, w \vDash \neg\varphi$ iff $\mathfrak{M}, w \nvDash \varphi$; and $\mathfrak{M}, w \vDash \varphi \wedge \psi$ iff $\mathfrak{M}, w \vDash \varphi$ and $\mathfrak{M}, w \vDash \psi$.

Lemma 5.4. Let $\mathfrak{M} = (W, V)$ be a possible worlds model and $\mathcal{M}$ the epistemic model of $(\wp(W), V)$. For any Boolean formula φ and (A, I) in $\mathcal{M}$, we have:

1. $\mathcal{M}, (A, I) \Vdash \varphi$ iff for all $w \in A$, we have $\mathfrak{M}, w \vDash \varphi$;

2. $\mathcal{M}, (A, I) \Vdash \Box\varphi$ iff for all $w \in I$, we have $\mathfrak{M}, w \vDash \varphi$;

3. $\mathcal{M}, (A, I) \Vdash \Diamond\varphi$ iff for some $w \in A$, we have $\mathfrak{M}, w \vDash \varphi$.

Proof. We prove part 1 by induction on φ. For the base case of a propositional variable $\mathtt{p}$, we have $\mathcal{M},(A,I)\Vdash \mathtt{p}$ iff $(A,I)\in V^{\mathsf{e}}(\mathtt{p})$ iff $A\subseteq V(\mathtt{p})$ iff for all $w\in A$, $\mathfrak{M},w\vDash \mathtt{p}$. The inductive step for $\wedge$ is immediate from the inductive hypothesis. For the inductive step for $\neg$, we have

$$\begin{aligned}\mathcal{M},(A,I)\Vdash\neg\varphi \quad &\text{iff}\quad \forall(A',I')\between(A,I),\ \mathcal{M},(A',I')\nVdash\varphi\\ &\text{iff}\quad \forall(A',I')\between(A,I)\ \exists w\in A':\mathfrak{M},w\nvDash\varphi \text{ by the inductive hypothesis}\\ &\text{iff}\quad \forall w\in A,\ \mathfrak{M},w\nvDash\varphi.\end{aligned}$$

The implication from the second to third line follows from the fact that for each $w\in A$, we have that $(\{w\},I)\between(A,I)$. The converse implication follows from the fact that $(A',I')\between(A,I)$ implies $A'\cap A\neq\varnothing$.

For part 2, we have

$$\begin{aligned}\mathcal{M},(A,I)\Vdash\Box\varphi \quad &\text{iff}\quad \forall(A',I')\in R((A,I)),\ \mathcal{M},(A',I')\Vdash\varphi\\ &\text{iff}\quad \forall(A',I')\in R((A,I))\ \forall w\in A',\ \mathfrak{M},w\vDash\varphi \text{ by part 1}\\ &\text{iff}\quad \forall w\in I,\ \mathfrak{M},w\vDash\varphi.\end{aligned}$$

The implication from the second to third line follows from the fact that $(A,I)R(I,I)$. The converse implication follows from the fact that $(A,I)R(A',I')$ implies $A'\subseteq I'\subseteq I$.

For part 3, we have

$$\begin{aligned}\mathcal{M},(A,I)\Vdash\Diamond\varphi \quad &\text{iff}\quad \forall(A',I')\between(A,I)\ \exists(A'',I'')\in R((A',I'))\\ &\qquad\qquad \exists(A''',I''')\between(A'',I''):\mathcal{M},(A''',I''')\Vdash\varphi\\ &\text{iff}\quad \forall(A',I')\between(A,I)\ \exists(A'',I'')\in R((A',I'))\\ &\qquad\qquad \exists(A''',I''')\between(A'',I'')\ \forall w\in A''',\mathfrak{M},w\vDash\varphi \text{ by part 1}\\ &\text{iff}\quad \exists w\in A:\mathfrak{M},w\vDash\varphi.\end{aligned}$$

For the implication from the second to third line, setting $(A',I')=(A,A)$, we have $(A',I')\between(A,I)$, and then from $(A',I')R(A'',I'')$ we have $A''\subseteq I''\subseteq I'\subseteq A$. Then from $(A''',I''')\between(A'',I'')$, we have $A'''\cap A''\neq\varnothing$ and hence $A'''\cap A\neq\varnothing$. Then since $\forall w\in A''',\mathfrak{M},w\vDash\varphi$, we have $\exists w\in A:\mathfrak{M},w\vDash\varphi$. For the converse implication, suppose $w\in A$, $\mathfrak{M},w\vDash\varphi$, and $(A',I')\between(A,I)$. It follows that $w\in I'$. Then take $(A'',I'')=(A'\cup\{w\},I')$, so $(A',I')R(A'',I'')$. Finally, taking $(A''',I''')=(\{w\},I')$, we have $(A''',I''')\between(A'',I'')$ and $\forall w\in A'''$, $\mathfrak{M},w\vDash\varphi$, which completes the proof. □

To see how Lemma 5.4 facilitates computations of truth values, let us consider a richer example starting from three worlds.

Example 5.5. Suppose we are about to roll a three-sided die, so $W=\{0,1,2\}$. Figure 16 shows the epistemic frame of $\wp(W)$. Figure 17 then shows the propositions expressed by several formulas, highlighted in green; the propositional variables $\mathtt{0}$, $\mathtt{1}$, and $\mathtt{2}$ are true in all possibilities (A,I) such that $A=\{0\}$, $A=\{1\}$, and $A=\{2\}$, respectively. Note the ease of computing the modal propositions using Lemma 5.4.

Finally, let us push one step further to the case of four worlds.

Example 5.6. Starting with $W=\{0,1,2,3\}$ is especially interesting, since we can capture all the possible truth value combinations for two Boolean propositional variables $\mathtt{p}$ and $\mathtt{q}$ with $V(\mathtt{p})=\{0,1\}$ and $V(\mathtt{q})=$

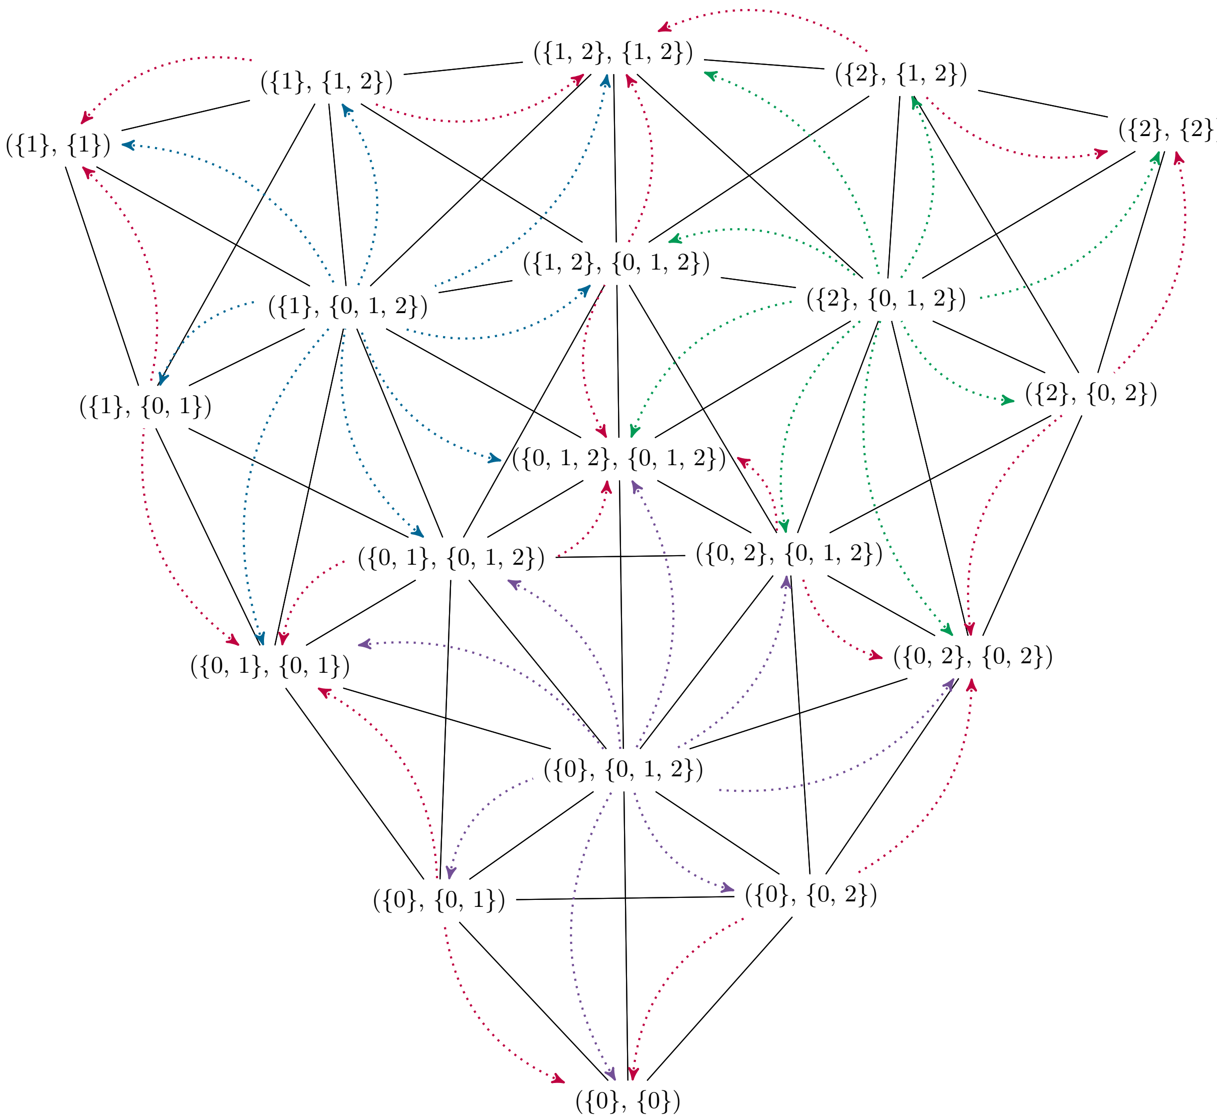

Figure 16: Epistemic frame constructed from three worlds

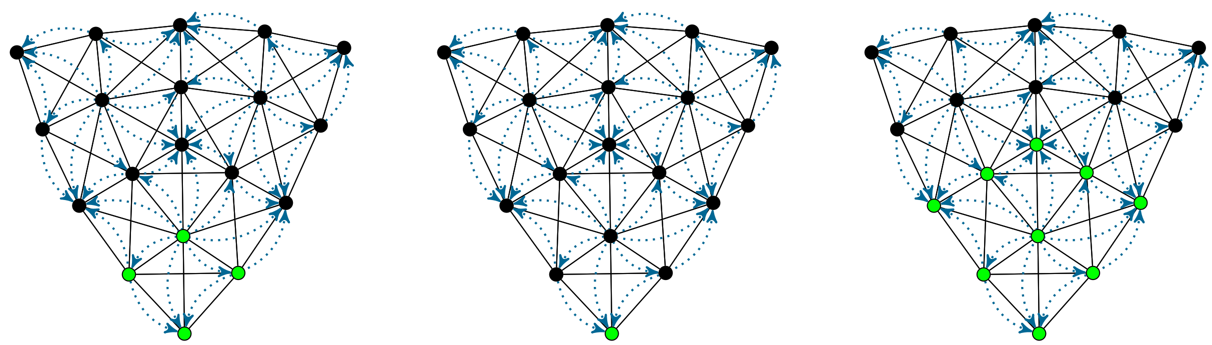

(a) $[\![\mathtt{0}]\!]^{\mathcal{M}}$ (b) $[\![\Box\mathtt{0}]\!]^{\mathcal{M}}$ (c) $[\![\Diamond\mathtt{0}]\!]^{\mathcal{M}}$

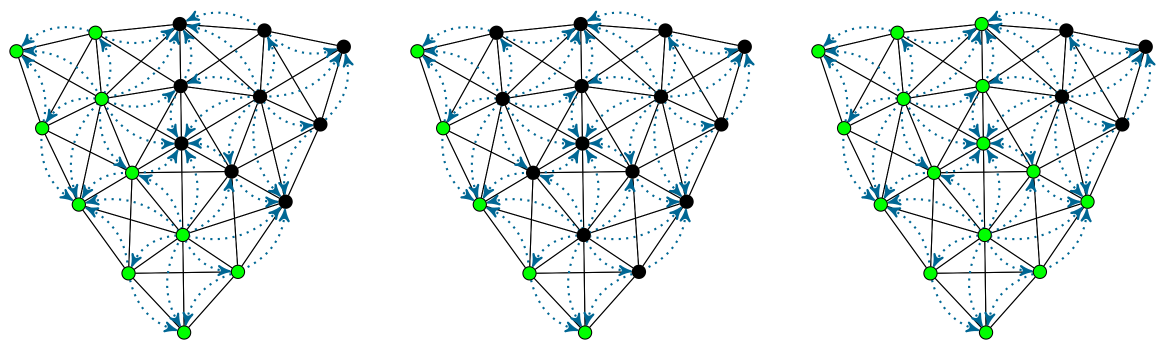

(d) $[\![\mathtt{0} \vee \mathtt{1}]\!]^{\mathcal{M}}$ (e) $[\![\Box(\mathtt{0} \vee \mathtt{1})]\!]^{\mathcal{M}}$ (f) $[\![\Diamond(\mathtt{0} \vee \mathtt{1})]\!]^{\mathcal{M}}$

Figure 17: Propositions highlighted in copies of the epistemic model constructed from the possible worlds model with $W = \{0, 1, 2\}$, $V(\mathtt{0}) = \{0\}$, $V(\mathtt{1}) = \{1\}$, and $V(2) = \{2\}$

$\{0, 3\}$. It is harder to cleanly draw the epistemic frame of $\wp(W)$ in this case, but Figure 18 provides a helpful three-dimensional tetrahedral visualization of the frame.

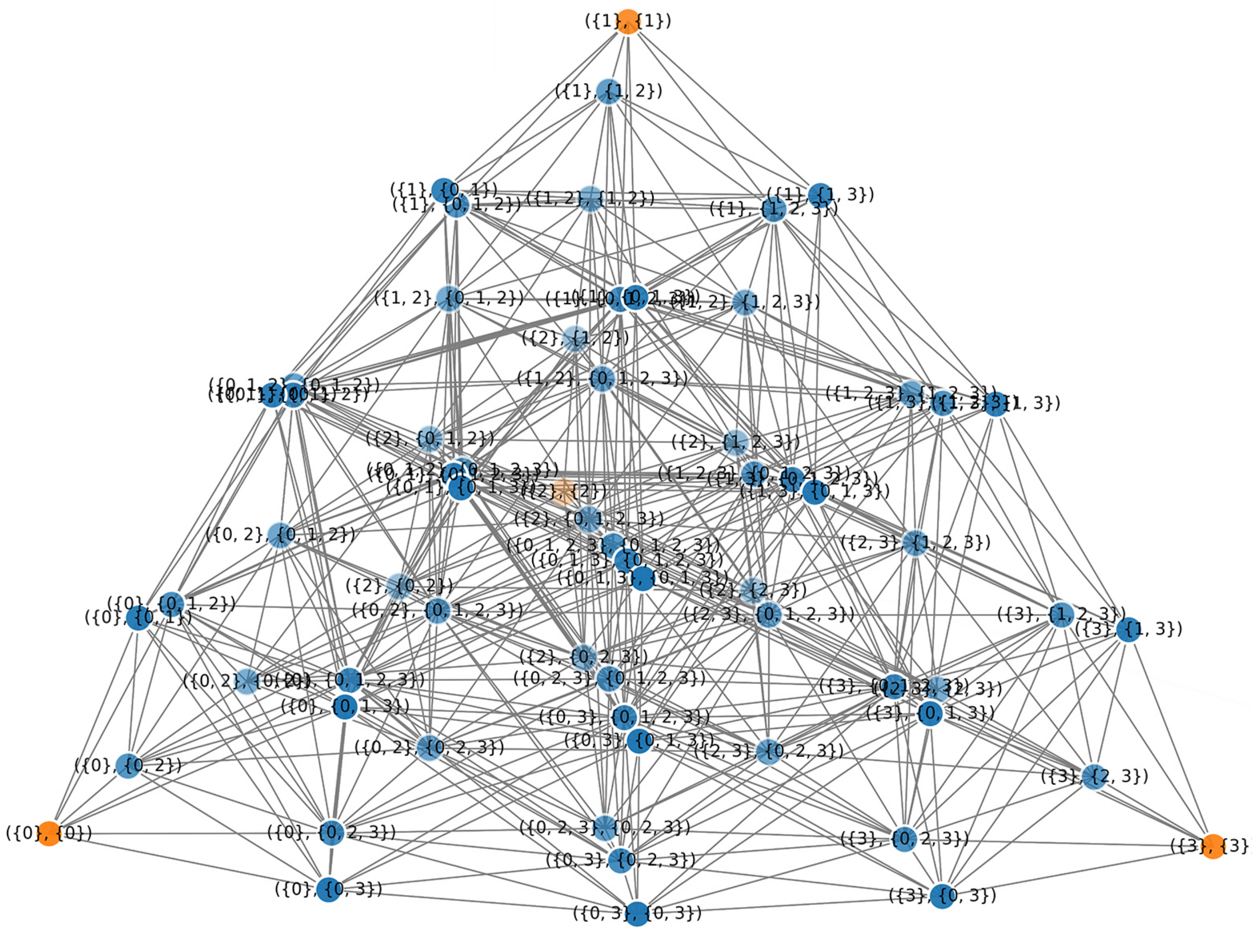

Figure 18: Epistemic frame constructed from four worlds with fully determinate possibilities highlighted. To reduce clutter, only compatibility relations are shown.

Table 1 shows the increase in the size of the epistemic frames as we increase the initial number of worlds. Visualization beyond four worlds becomes difficult, but our computer implementation of the construction can comfortably handle more than four initial worlds (though calculating the size of the ortholattice coming from the epistemic frame becomes prohibitive without an analytic expression for this size).

Having seen the epistemic frames constructed from three Boolean algebras, we now prove the first main result of this section. We show that the epistemic frame of B is an epistemic compatibility frame (Definition 4.26) into whose associated epistemic ortholattice B embeds. The bare fact that B embeds into an (S5) epistemic ortholattice is trivial, since B itself equipped with $\Box$ as the identity function—and hence where $\Diamond$ is the identity function—is an (S5) epistemic ortholattice. But what we want is to embed B into an epistemic ortholattice where $\Diamond$ does not collapse, as in part 4 of the following.

Theorem 5.7. For any Boolean algebra B with lattice order $\leq$:

1. B^{e} is an epistemic compatibility frame;

size of B	world representation of B	epistemic frame of B	epistemic extension of B
4	2	5	10
8	3	19	302
16	4	65	298,414
32	5	211	?
64	6	665	?

Table 1: Sizes of objects associated with a Boolean algebra B.

2. the map e defined by $e_B(a) = \{(b,i) \in S \mid b \leq a\}$ is an embedding of B into the complete epistemic ortholattice $O(B^{\mathsf{e}})$, which we therefore call the *epistemic extension of* B;

3. $O(B^{\mathsf{e}})$ is an S5 epistemic ortholattice;

4. for all $b \in B$, if $b \notin \{0,1\}$, then $\Diamond e_B(b) \not\leq e_B(b)$ in $O(B^{\mathsf{e}})$.

Proof. For part 1, clearly $\between$ is reflexive and symmetric, so we proceed to the properties of R.

For R-regularity in its form in Lemma 4.21, suppose $(a,i)R(b,k) \between (d,l)$. We claim that the desired witness for R-regularity is $(a',i') = (a \vee d, a \vee d)$. First, we claim $(a,i) \between (a \vee d, a \vee d)$. The first two conditions, namely $a \wedge (a \vee d) \neq 0$ and $a \leq a \vee d$, are immediate. As for $a \vee d \subseteq i$, from $(a,i) \in S$ we have $a \leq i$, and from $(a,i)R(b,k) \between (d,l)$ we have $i \geq k \geq d$. Now suppose $(a'',i'') \between (a \vee d, a \vee d)$. Then we claim $(a'',i'')R(a \vee d, a \vee d) \between (d,l)$. That $(a'',i'')R(a \vee d, a \vee d)$ is immediate from $(a'',i'') \between (a \vee d, a \vee d)$. Finally, from $(d,l) \in S$ we have $d \leq l$, and from $(a,i)R(b,k) \between (d,l)$ we have $a \leq b \leq l$, so $a \leq d \leq l$. Hence $(a \vee d, a \vee d) \between (d,l)$. This completes the proof of R-regularity.

For Reflexivity, it is immediate from the definition of R that $(a,i)R(a,i)$.

For Knowability, given $x = (a,i)$, we claim that $y = (a,a)$ is the desired witness. For suppose $z = (a',i')$ and (i) $(a,a)R(a',i')$. We claim that z refines x. Supposing (ii) $(a'',i'') \between (a',i')$, we must show that $(a'',i'') \between (a,i)$, i.e., that (a) $a'' \wedge a \neq 0$, (b) $a'' \leq i$, and (c) $a \leq i''$. By (i), we have $a' \leq i' \leq a$, and by (ii), we have $a'' \wedge a' \neq 0$, so (a) follows. By (ii) we also have $a'' \leq i'$, which with $i' \leq a$ from (i) and $a \leq i$ implies (b). Finally, by (ii) we have $a' \leq i''$, and by (i) we have $a \leq a'$, so (c) follows.

For part 2, we first prove that $e_B(a)$ is $\between$-regular. Suppose $(b,i) \notin e_B(a)$, so $b \not\leq a$. Then where $b' = b \wedge \neg a$, we have $0 \neq b' \leq b \leq i$, so $(b',i) \between (b,i)$. Now consider any $(b'',i'') \between (b',i)$. Then $b'' \wedge b' \neq 0$, which implies $b'' \not\leq a$, so $(b'',i'') \notin e_B(a)$. Thus, we have shown that if $(b,i) \notin e_B(a)$, then there is a $(b',i') \between (b,i)$ such that for all $(b'',i'') \between (b',i')$, we have $(b'',i'') \notin e_B(a)$. Hence $e_B(a)$ is $\between$-regular.

Next, clearly e preserves all existing meets from B:

$$e_B(\bigwedge_{a \in A} a) = \{(b,i) \in S \mid b \leq \bigwedge_{a \in A} a\} = \bigcap_{a \in A} \{(b,i) \in S \mid b \leq a\} = \bigwedge_{a \in A} e_B(a).$$

For preservation of joins, we need only show that $e_B(\bigvee A) \subseteq \bigvee\{e_B(a) \mid a \in A\}$, as the converse follows from order preservation, which follows from meet preservation. Suppose $(b,i) \in e_B(\bigvee A)$, so $b \leq \bigvee A$, and $(b',i') \between (b,i)$, so $b' \wedge b \neq 0$. Since $0 \neq b' \wedge b \leq \bigvee A$, by distributivity in B there is some $a \in A$ such that $a \wedge b' \wedge b \neq 0$. Then setting $b'' = a \wedge b' \wedge b$, we have $(b'',i') \between (b,i)$ and $(b'',i') \in e_B(a)$. This shows that $(b,i) \in \bigvee\{e_B(a) \mid a \in A\}$.

Finally, we show that $e_B(\neg a) = \neg e_B(a)$. Suppose $(b,i) \in e_B(\neg a)$, so $b \leq \neg a$. Then for any $(b',i') \between (b,i)$, since $b' \wedge b \neq 0$, it follows that $b' \not\leq a$, so $(b',i') \notin e_B(a)$, which shows that $(b,i) \in \neg e_B(a)$. Conversely, suppose $(b,i) \notin e_B(\neg a)$, so $b \not\leq \neg a$. Then $(b \wedge a, i) \between (b,i)$ and $(b \wedge a, i) \in e_B(a)$, so $(b,i) \notin \neg e_B(a)$.

For part 3, for all $U \in O(B^{\mathsf{e}})$, that $\Box U \subseteq U$ is immediate from the reflexivity of the accessibility relation R, and that $\Box U \subseteq \Box\Box U$ is immediate from the transitivity of R. From here it suffices to show $U \subseteq \Box\Diamond U$ for all $U \in O(B^{\mathsf{e}})$. Suppose $(A, I) \in U$ and $(A, I)R(A', I')$, so $A \subseteq A'$ and $I' \subseteq I$. To show $(A', I') \in \Diamond U$, suppose $(A'', I'') \between (A', I')$, so $A'' \subseteq I'$ and $A' \subseteq I''$. From $A \subseteq A' \subseteq I''$ and $A'' \subseteq I''$, we have $(A'', I'')R(A \cup A'', A \cup A'')$, and from $A \subseteq I$ and $A'' \subseteq I' \subseteq I$, we have $(A \cup A'', A \cup A'') \between (A, I) \in U$. This shows that $(A, I) \in \Box\Diamond U$.

For part 4, if $b \neq 0$, then it is easy to see that $(1, 1) \in \Diamond e_B(b)$, and if $b \neq 1$, then $(1, 1) \notin e_B(b)$. $\Box$

Our next goal is to show that $O(B^{\mathsf{e}})$ equipped with the distinguished Boolean subalgebra $\mathbb{B} = \{e_B(b) \mid b \in B$ is a complete epistemic *ortho-Boolean lattice* as in Definition 3.32, or equivalently, that B^{e} equipped with the distinguished $\mathbb{B} = \{e_B(b) \mid b \in B\}$ is a *stratified* epistemic compatibility frame as in Definition 4.40. This is equivalent to showing that each $\mathbb{B}_n$ in the hierarchy of Definition 4.40 is a Boolean algebra. To prove this, we first prove a useful algebraic generalization of the main idea of Lemma 5.4.

Lemma 5.8. For any Boolean algebra B, the following are equivalent for any possibility (a, i) in B^{e} and families $\{b_k, d_k^1, \dots, d_k^{m_k}\}$ of elements of B for $k \in \{1, \dots, n\}$:

1. $(a, i) \in \bigvee\limits_{1 \leq k \leq n} \big(\Box e_B(b_k) \wedge \Diamond e_B(d_k^1) \wedge \dots \wedge \Diamond e_B(d_k^{m_k})\big)$;
2. for any c such that $a \leq c \leq i$, there is some k such that $c \leq b_k$ and $c \wedge d_k^\ell \neq 0$ for all $\ell \in \{1, \dots, m_k\}$.

Proof. By essentially the same reasoning used in the proof of Lemma 5.4, only with propositions instead of formulas and an arbitrary Boolean algebra instead of a powerset, we have that for any (a, i) in the epistemic frame, $(a, i) \in \Box e_B(b)$ iff $i \leq b$, and $(a, i) \in \Diamond e_B(d)$ iff $a \wedge d \neq 0$. Using these equivalences, we now prove the more general statement above.

Suppose part 1 holds and $a \leq c \leq i$. Then $(c, c) \between (a, i)$, so there is an $(a', i') \between (c, c)$ and $k \in \{1, \dots, n\}$ such that $(a', i') \in \Box e_B(b_k) \wedge \Diamond e_B(d_k^1) \wedge \dots \wedge \Diamond e_B(d_k^{m_k})$. Then $i' \leq b_k$ and $a' \wedge d_k^\ell \neq 0$ for all $\ell \in \{1, \dots, m_k\}$. Then since $(a', i') \between (c, c)$ implies $a' \leq c$ and $c \leq i'$, it follows that $c \leq b_k$ and $c \wedge d_k^\ell \neq 0$ for all $\ell \in \{1, \dots, m_k\}$. Therefore, part 2 holds.

Now suppose part 2 holds. Suppose $(a', i') \between (a, i)$. Then since $a \leq a \vee a' \leq i$, part 2 implies there is some k such that $a \vee a' \leq b_k$ and $(a \vee a') \wedge d_k^\ell \neq 0$ for all $\ell \in \{1, \dots, m_k\}$. From $(a', i') \between (a, i)$ we also have $a \vee a' \leq i'$. Then where

$$x = a' \vee \bigvee_{1 \leq \ell \leq m_k} ((a \vee a') \wedge d_k^\ell),$$

we have $x \wedge a' \neq 0$, $x \leq i'$, and $a' \leq b_k$, so $(x, b_k) \between (a', i')$, and $(x, b_k) \in \Box e_B(b_k) \wedge \Diamond e_B(d_k^1) \wedge \dots \wedge \Diamond e_B(d_k^{m_k})$. Thus, we have shown that for every $(a', i') \between (a, i)$, there is an $(a'', i'') \between (a', i')$ and $k \in \{1, \dots, m_k\}$ such that $(a'', i'') \in \Box e_B(b_k) \wedge \Diamond e_B(d_k^1) \wedge \dots \wedge \Diamond e_B(d_k^{m_k})$. Therefore, part 1 holds. $\Box$

Using Lemma 5.8, we can show that a general distributive law holds for propositions resulting from applying $\Box$ and $\Diamond$ to Boolean propositions.

Lemma 5.9. Let B be a Boolean algebra. Let $\{U_{t,j} \mid t \in T, j \in J_t\}$ be a finite family of propositions in $O(B^{\mathsf{e}})$ such that for each $U_{t,j}$, there is a $b \in B$ such that $U_{t,j} = \Box e_B(b)$ or $U_{t,j} = \Diamond e_B(b)$. Then we have

$$\bigwedge_{t \in T} \bigvee_{j \in J_t} U_{t,j} \subseteq \bigvee_{f \in F} \bigwedge_{t \in T} U_{t, f(t)}$$

where F is the set of functions that choose for each $t \in T$ some $f(t) \in J_t$.

Proof. Suppose $(a,i) \in \bigwedge_{t\in T} \bigvee_{j\in J_t} U_{t,j}$. Then by Lemma 5.8, for each $t \in T$ and c such that $a \leq c \leq i$, there is some $j \in J_t$ such that either

(i) $U_{t,j} = \Box e_B(b)$ and $c \leq b$ or

(ii) $U_{t,j} = \Diamond e_B(d)$ and $c \wedge d \neq 0$.

Let f_c be a function that chooses such $j \in J_t$ for each $t \in T$. Then let us write $\bigwedge_{t\in T} U_{t,f_c(t)}$ as

$$\Box e_B(b_c^1) \wedge \cdots \wedge \Box e_B(b_c^{n_c}) \wedge \Diamond e_B(d_c^1) \wedge \cdots \wedge \Diamond e_B(d_c^{m_c}).$$

Where $b_c = b_c^1 \wedge \cdots \wedge b_c^{n_c}$, this is equivalent to

$$\Box e_B(b_c) \wedge \Diamond e_B(d_c^1) \wedge \cdots \wedge \Diamond e_B(d_c^{m_c}).$$

Then to show that $(a,i) \in \bigvee_{f\in F} \bigwedge_{t\in T} U_{t,f(t)}$, it suffices by Lemma 5.8 to show that for each c such that $a \leq c \leq i$, we have $c \leq b_c$ and $c \wedge d_c^s \neq 0$ for each $s \in \{1,\ldots,m_c\}$. Indeed, we have $c \leq b_c$ by (i) in the construction of f_c, and $c \wedge d_c^s \neq 0$ for each $s \in \{1,\ldots,m_c\}$ by (ii) in the construction of f_c. This completes the proof. ☐

We are now ready to prove that each of the distinguished algebras $\mathbb{B}_n$ is indeed Boolean.

Theorem 5.10. For any Boolean algebra B, $O(B^{\mathsf{e}})$ with the distinguished Boolean subalgebra $\{e_B(b) \mid b \in B\}$ is a complete epistemic ortho-Boolean lattice. Equivalently, $(B^{\mathsf{e}}, \{e_B(b) \mid b \in B\})$ is a stratified epistemic compatibility frame.

Proof. Where $\mathbb{B} = \{e_B(b) \mid b \in B\}$, it is immediate from Theorem 5.7.2 that $\mathbb{B}$ forms a Boolean algebra. It only remains to show that each $\mathbb{B}_n$ for $n \geq 1$ is a Boolean algebra. First, let $\mathbb{B}_1'$ be the set of elements generated from $\{\Box U \mid U \in \mathbb{B}_0\} \cup \{\Diamond U \mid U \in \mathbb{B}_0\}$ by $\wedge$ and $\vee$. Thanks to De Morgan's laws and involution for negation, $\mathbb{B}_1' = \mathbb{B}_1$. Now we claim that each $V \in \mathbb{B}_1'$ is equal to $\Box V$. The proof is by induction, given the inductive definition of $\mathbb{B}_1'$. In the case where V is $\Box U$ or $\neg\Box U$ for some $U \in \mathbb{B}_0$, we have that $\Box U = \Box\Box U$ and $\neg\Box U = \Box\neg\Box U$ by the S5 axioms given by Theorem 5.7.3. If V is $V_1 \wedge V_2$ for $V_1, V_2 \in \mathbb{B}_1'$, then by the inductive hypothesis, $V_1 \wedge V_2 = \Box V_1 \wedge \Box V_2$, which is equal to $\Box(V_1 \wedge V_2)$. Finally, if V is $V_1 \vee V_2$ for $V_1, V_2 \in \mathbb{B}_1'$, then by the inductive hypothesis, $V_1 \vee V_2 = \Box V_1 \vee \Box V_2$, which is equal (using the 4 and T axioms) to $\Box(\Box V_1 \vee \Box V_2)$. Thus, each $V \in \mathbb{B}_1$ is equal to $\Box V$. Hence each generator of $\mathbb{B}_2$, i.e., each $\Box V$ for $V \in \mathbb{B}_1$, already belongs to $\mathbb{B}_1$, which implies $\mathbb{B}_2 = \mathbb{B}_1$ and indeed $\mathbb{B}_n = \mathbb{B}_1$ for $n \geq 1$. Thus, to complete the proof we need only show that $\mathbb{B}_1$, or equivalently, $\mathbb{B}_1'$, is distributive. As shown by Kolibiar (1972, Thm. 2a), to show that $\mathbb{B}_1'$ is distributive and hence Boolean, it suffices to show that for any finite set M of finite subsets of the generating set $\{\Box U \mid U \in \mathbb{B}_0\} \cup \{\Diamond U \mid U \in \mathbb{B}_0\}$, we have

$$\bigwedge_{X\in M} \bigvee X \subseteq \bigvee_{t\in\Pi M} \bigwedge_{X\in M} t(X),$$

where ΠM is the set of all functions t on M such that $t(X) \in X$ for $X \in M$. This is precisely Lemma 5.9. ☐

The moral of Theorems 5.7 and 5.10 is that we can add to any Boolean algebra B of propositions new epistemic modal propositions, by embedding B into a complete epistemic ortho-Boolean lattice that arises from a stratified epistemic compatibility frame built from B. The resulting "epistemic extension" therefore validates our logic EO^+. Thus, we have an abundant Boolean source of models for our logic.

Now the obvious questions are the following. Which complete epistemic ortho-Boolean lattices arise as epistemic extensions of Boolean algebras? And what is the logic of epistemic frames of Boolean algebras? We answer the first question with an algebraic characterization in Appendix A.1. In Appendix A.2, we prove the *decidability* of the logic of epistemic frames of Boolean algebras; and in Appendix A.3, we prove a related semantic normal form theorem. Crucially, the logic is stronger than EO^+, not only because it includes the S5 principles but also because it includes the following principles.

Proposition 5.11. For any Boolean algebra B and $U, U_j, V_{j,k} \in O(B^{\mathsf{e}})$ in the image of the embedding e_B:

1. restricted diamond distributivity: if $V_{j,k} \subseteq U_j$ for all $1 \leq j \leq n$ and $1 \leq k \leq m_j$, then

$$(\bigvee_{1\leq j\leq n} U_j) \wedge \bigwedge_{1\leq j\leq n} (\Diamond V_{j,1} \wedge \cdots \wedge \Diamond V_{j,m_j}) \subseteq \bigvee_{1\leq j\leq n} (U_j \wedge \Diamond V_{j,1} \wedge \cdots \wedge \Diamond V_{j,m_j});$$

2. box distributivity: $(U_1 \vee U_2) \wedge \Box V \subseteq (U_1 \wedge \Box V) \vee (U_2 \wedge \Box V)$;

3. inheritance: $(U \wedge \Diamond V) \subseteq \Diamond(U \wedge V)$;

4. diamond disjunctive syllogism A: $(U \vee \Diamond V) \wedge \neg\Diamond V \subseteq U$;

5. diamond disjunctive syllogism B: $(U \vee \Diamond V) \wedge \neg U \subseteq \Diamond V$.

Proof. For part 1, let $b_j = e_B^{-1}(U_j)$ and $c_{j,k} = e_B^{-1}(V_{j,k})$. Then since $V_{j,k} \subseteq U_j$, we have $c_{j,k} \leq b_j$. Suppose (a,i) is in the left-hand side of the inclusion. Then by Lemma 5.8, we have $a \leq b_1 \vee \cdots \vee b_n$ and $a \wedge c_{j,k} \neq 0$ for $1 \leq j \leq n$ and $1 \leq k \leq m_j$. To show (a,i) is in the right-hand side, consider any $(a',i') \between (a,i)$. Since $a' \wedge a \neq 0$ and $a \leq b_1 \vee \cdots \vee b_n$, it follows that $a' \wedge a \wedge b_j \neq 0$ for some j. Let $x = (a' \wedge a \wedge b_j) \vee (a \wedge c_{j,1}) \vee \cdots \vee (a \wedge c_{j,m_j})$. Since $a' \leq i'$ and $a \leq i'$, we have $x \leq i'$. Hence $(x,i') \between (a',i')$. Moreover, since $x \leq b_j$ and $x \wedge c_{j,k} \neq 0$ for $1 \leq k \leq m_j$, we have $(x,i') \in U_j \wedge \Diamond V_{j,1} \wedge \cdots \wedge \Diamond V_{j,m_j}$ by Lemma 5.8. Thus, for every $(a',i') \between (a,i)$, there is an $(a'',i'') \in (U_1 \wedge \Diamond V_{1,1} \wedge \cdots \wedge \Diamond V_{1,m_1}) \cup \cdots \cup (U_n \wedge \Diamond V_{n,1} \wedge \cdots \wedge \Diamond V_{n,m_n})$, which shows that (a,i) is in the right-hand side.

For part 2, let $b_1 = e_B^{-1}(U_1)$, $b_2 = e_B^{-1}(U_2)$, and $c = e_B^{-1}(V)$. Suppose (a,i) is in the left-hand side, so $i \leq c$ by Lemma 5.8. Further suppose that $(a',i') \between (a,i)$, so $a' \leq i \leq c$. Then since $(a,i) \in U_1 \vee U_2$, there is an $(a'',i'') \between (a',i')$ such that $(a'',i'') \in U_i$ for some $i \in \{1,2\}$, so $a'' \leq b_i$ by Lemma 5.8. Now $(a'',i'') \between (a',i')$ implies $(a'' \wedge a', a') \between (a',i')$, and since $a'' \leq b_i$ and $a' \leq c$, we have $(a'' \wedge a', a') \in U_i \wedge \Box V$ by Lemma 5.8. Thus, for every $(a',i') \between (a,i)$, there is an $(a'',i'') \in (U_1 \wedge \Box V) \cup (U_2 \wedge \Box V)$, which shows that (a,i) is in the right-hand side.

For part 3, let $b = e_B^{-1}(U)$ and $c = e_B^{-1}(V)$. If (a,i) is in the left-hand side of the inclusion, then by Lemma 5.8, $a \leq b$ and $a \wedge c \neq 0$, which implies $a \wedge b \wedge c \neq 0$ and hence (a,i) is in the right-hand side.

For parts 4 and 5, let $b = e_B^{-1}(U)$ and $c = e_B^{-1}(V)$. If $(a,i) \in U \vee \Diamond V$, then for any $(a',i') \between (a,i)$ there is an $(a'',i'') \between (a',i')$ such that $(a'',i'') \in U$ or $(a'',i'') \in \Diamond V$, so $a'' \leq b$ or $a'' \wedge c \neq 0$ by Lemma 5.8.

For part 4, if $(a,i) \notin U$, so $a \not\leq b$ by Lemma 5.8, then $(a \wedge \neg b, i) \between (a,i)$, so by the previous paragraph, there is an $(a'',i'') \between (a \wedge \neg b, i)$ such that $a'' \leq b$ or $a'' \wedge c \neq 0$; but $a'' \leq b$ contradicts $a'' \wedge a \wedge \neg b \neq 0$ from $(a'',i'') \between (a \wedge \neg b, i)$, so we conclude that $a'' \wedge c \neq 0$, which implies $(a'', i'' \vee a) \in \Diamond V$ by Lemma 5.8. Finally, $(a'',i'') \between (a \wedge \neg b, i)$ implies $(a'', i'' \vee a) \between (a,i)$, so $(a'', i'' \vee a) \in \Diamond V$ implies $(a,i) \notin \neg\Diamond V$.

For part 5, if $(a,i) \notin \Diamond V$, so $a \wedge c = 0$ by Lemma 5.8, then let $(a',i') = (a,a)$. Hence $(a',i') \between (a,i)$, so as above there is an $(a'',i'') \between (a',i')$ such that $a'' \leq b$ or $a'' \wedge c \neq 0$. But $(a'',i'') \between (a',i')$ implies $a'' \leq i' = a$, so $a \wedge c = 0$ implies $a'' \wedge c = 0$ and hence $a'' \leq b$, which implies $(a'',i'') \in U$. Since $(a'',i'') \between (a',i') = (a,a)$, it follows that $(a'',i'') \between (a,i)$, so $(a,i) \notin \neg U$. $\square$

It is not difficult to construct epistemic ortho-Boolean lattices validating EO^+ that invalidate the principles in Proposition 5.11. These principles may, however, by desirable (though the plausibility of part 3 may depend on the S5 assumption built into the construction of possibilities from a set of worlds). To see the plausibility of the first principle, consider examples like (24), suggested to us by Seth Yalcin (p.c.):

(24) a. Either the window is open or the window is closed, and it might be raining in France.
b. Therefore, either the window is open and it might be raining in France, or the window is closed and it might be raining in France.

The inference from (24-a) to (24-b) is intuitively reasonable. After all, the possible weather in France has nothing to do with the state of my window, and so nothing, intuitively, prevents us from distributing the disjunction over the conjunction in this case. Given the enthymematic premise that *it might be that both the window is open and it's raining in France*, and *it might be that both the window is closed and it's raining in France*, the weakening of distributivity in Proposition 5.11.1 would explain the felt validity of this reasoning. The other principles also strike us as reasonable, though we will not discuss them further here; see Goldstein 2019; Hawke and Steinert-Threlkeld 2021 for discussion of inheritance in particular.

The principles in Proposition 5.11 highlight the interest of providing an axiomatization of the logic of epistemic frames of Boolean algebras, extending EO^+, in the style of Definitions 3.24 and 3.38. However, we leave this task for future work.

5.2 Lifting probabilities

In this section, we use the epistemic extension construction to accomplish a partial reconciliation between classical probability and the non-classical behavior of epistemic modals.

The immediate problem is that it seems perfectly reasonable to assign high probability to p and high probability, even probability 1, to $\Diamond\neg p$, despite the fact that $p \wedge \Diamond\neg p$ is inconsistent and hence should never have non-zero probability. Yet this is inconsistent with classical probability. Abstractly, this puzzle is the measure-theoretic corollary of the failure of pseudocomplementation for epistemic modals. The corresponding measure-theoretic principle says that, if $\mu(p\wedge q) = 0$, then $\mu(p) \leq 1-\mu(q)$. This is a law of classical probability theory, but it appears to fail for epistemic modals: intuitively, $p \wedge \Diamond\neg p$ always has probability 0, while the probability of p and the probability of $\Diamond\neg p$ can sum to more than 1. More concretely, suppose that a reliable weather report tells you there is a high likelihood of rain later, say .9, and this is all the information you have about the weather. Then you should intuitively have credence .9 that it will rain. But you should also be sure that it might not rain: that is, you should have credence around 1 that it might not rain. If your credences are probabilistic, then it follows from the laws of probability that you should have credence around .9 that it will rain and might not rain. But that is intuitively wrong: just as it seems impossible to imagine that it will rain and might not, it likewise seems impossible to reasonably assign any credence to the proposition that it will rain and might not.

Hence the probabilistic analogue of pseudocomplementation fails for epistemic modals. Since our possibility semantics already yields a compelling model of the failure of pseudocomplementation, it is natural to look for a way to extend a classical probability measure to a probability function on a set of modal propositions that accounts for this non-classicality. And indeed, the epistemic extension construction suggests just such an extension of probabilities, as follows.[29]

[29]For simplicity, we only consider probability measures on the full powerset of the set of worlds.

Definition 5.12. Given a nonempty set W, distinguished information state $\mathcal{I} \subseteq W$, and a finitely additive probability measure $\mu : \wp(W) \to [0,1]$ with $\mu(\mathcal{I}) = 1$, we define the *epistemic extension* $\mu^{\mathsf{e}}_{\mathcal{I}} : O(\wp(W)^{\mathsf{e}}) \to [0,1]$ of μ with respect to U as follows:

- $\mu^{\mathsf{e}}_{\mathcal{I}}(U) = \mu\big(\bigcup\{A \subseteq W \mid (A, \mathcal{I}) \in U\}\big)$.

A natural choice of $\mathcal{I}$, at least in the finite case, is $\mathcal{I} = \{w \in W \mid \mu(\{w\}) > 0\}$.

Intuitively, to compute the probability of a proposition $U \in O(\wp(W)^{\mathsf{e}})$, we compute the probability of the worldly proposition obtained by unioning the first coordinates of those possibilities $(A, \mathcal{I}) \in U$. A useful fact is that this union is either empty or yields the *largest* A such that $(A, \mathcal{I}) \in U$, thanks to the following.

Lemma 5.13. For any Boolean algebra B, possibilities (a, i) and (b, i) in B^{e} and proposition $U \in O(B^{\mathsf{e}})$, if $(a, i) \in U$ and $(b, i) \in U$, then $(a \vee b, i) \in U$.

Proof. Suppose $(a, i) \in U$ and $(b, i) \in U$. Further suppose $(a', i') \between (a \vee b, i)$, so $a' \wedge (a \vee b) \neq 0$, $a' \leq i$, and $a \vee b \leq i'$. Since $a' \wedge (a \vee b) \neq 0$, we have either $a' \wedge a \neq 0$ or $a' \wedge b \neq 0$. Suppose $a' \wedge a \neq 0$. Then $(a', i') \between (a, i)$, which with $(a, i) \in U$ and the $\between$-regularity of U implies there is an $(a'', i'') \between (a', i')$ such that $(a'', i'') \in U$. The argument if $a' \wedge b \neq 0$ is analogous with (b, i). Thus, for every $(a', i') \between (a \vee b, i)$, there is an $(a'', i'') \between (a', i')$ such that $(a'', i'') \in U$, which shows $(a \vee b, i) \in U$ by the $\between$-regularity of U. $\square$

Example 5.14. In the context of Example 5.5 where $W = U = \{0, 1, 2\}$ and μ is the uniform measure with $\mu(\{0\}) = \mu(\{1\}) = \mu(\{2\}) = 1/3$, we obtain the lifted probabilities in Table 2. For example, we calculate that

$$\begin{aligned}
\mu^{\mathsf{e}}_{W}(\llbracket \Box \mathsf{0} \rrbracket^{\mathcal{M}}) &= \mu\Big(\bigcup\big\{A \subseteq W \mid (A, W) \in \llbracket \Box \mathsf{0} \rrbracket^{\mathcal{M}})\big\}\Big) \\
&= \mu\Big(\bigcup\big\{A \subseteq W \mid (A, W) \in \{(\{0\}, \{0\})\}\big\}\Big) \\
&= \mu(\varnothing) = 0.
\end{aligned}$$

formula φ	$\mu^{\mathsf{e}}_{W}(\llbracket \varphi \rrbracket^{\mathcal{M}})$
$\mathsf{0}$	$1/3$
$\Box \mathsf{0}$	0
$\Diamond \mathsf{0} \wedge \Diamond \mathsf{1} \wedge \Diamond \mathsf{2}$	1
$\mathsf{0} \wedge \Diamond \mathsf{1}$	0

Table 2: Lifted probabilities given the uniform distribution on three worlds

Example 5.15. Figure 19 shows the probabilities of all propositions in the epistemic frame from two worlds as in Example 5.3 when we start with the uniform measure with $\mu(\{0\}) = \mu(\{1\}) = 1/2$.

The following fact shows that classical probability is fully preserved for non-modal propositions.

Proposition 5.16. For any W, $\mathcal{I}$, and μ as in Definition 5.12, and proposition $C \subseteq W$,

$$\mu^{\mathsf{e}}_{\mathcal{I}}(e(C)) = \mu(C)$$

where e is the embedding of $\wp(W)$ into $O(\wp(W)^{\mathsf{e}})$ from Theorem 5.7.2.

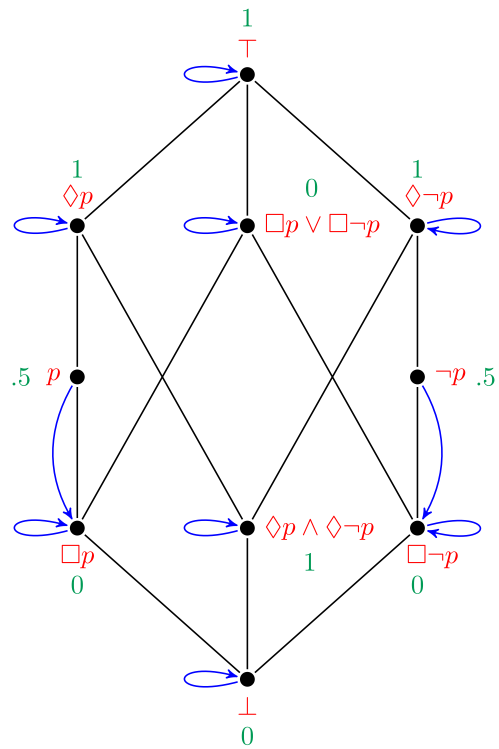

Figure 19: Probability values (green) on the ortholattice of the Epistemic Scale from Example 5.3.

Proof. We have

$$\begin{aligned} \mu^{\mathsf{e}}_{\mathcal{I}}(e(C)) &= \mu\big(\bigcup\{A \subseteq W \mid (A, \mathcal{I}) \in e(C)\}\big) \text{ by definition of } \mu^{\mathsf{e}}_{\mathcal{I}} \\ &= \mu\big(\bigcup\{A \subseteq W \mid A \subseteq \mathcal{I} \text{ and } A \subseteq C\}\big) \text{ by definition of } e \\ &= \mu(\mathcal{I} \cap C) \\ &= \mu(C) \text{ since } \mu(\mathcal{I}) = 1. \quad \square \end{aligned}$$

Corollary 5.17. Where $\mathbb{B}_0$ is image of the embedding e, the function $\mu^{\mathsf{e}}_{\mathcal{I}}$ restricted to $\mathbb{B}_0$ is a finitely additive probability measure.

Turning to modal propositions, the key observation is that the epistemic extension of a probability measure will capture exactly the failures of the measure-theoretic corollary of pseudocomplementation that we want to model. For instance, in Figure 19, $p \wedge \Diamond\neg p$ has probability 0, but the probability of p and the probability of $\Diamond\neg p$ sum to 1.5. More generally, the epistemic extension of a probability measure can yield probability functions that assign arbitrarily high probability to p short of 1 and arbitrarily high probability to $\Diamond\neg p$; but these probability functions will always assign probability 0 to $p \wedge \Diamond\neg p$, hence capturing the non-classical probabilistic behavior of epistemic modals.

However, other properties of classical probability do hold for arbitrary propositions.

Proposition 5.18. For any W, $\mathcal{I}$, and μ as in Definition 5.12 and $U, V \in O(\wp(W)^{\mathsf{e}})$:

1. if $U \subseteq V$, then $\mu^{\mathsf{e}}_{\mathcal{I}}(U) \leq \mu^{\mathsf{e}}_{\mathcal{I}}(V)$;

2. $\mu^{\mathsf{e}}_{\mathcal{I}}(\neg U) = 1 - \mu^{\mathsf{e}}_{\mathcal{I}}(U)$.

Proof. Part 1 is immediate from Definition 5.12 and the finite additivity of μ.

For part 2, we prove the equivalent $\mu^{\mathsf{e}}_{\mathcal{I}}(U) = 1 - \mu^{\mathsf{e}}_{\mathcal{I}}(\neg U)$.

Case 1: $(\mathcal{I},\mathcal{I}) \in U$, so $\mu^{\mathsf{e}}_{\mathcal{I}}(U) = 1$ by Definition 5.12. Then given $(A,\mathcal{I}) \between (\mathcal{I},\mathcal{I})$ for any $A \subseteq \mathcal{I}$, there is no A with $(A,\mathcal{I}) \in \neg U$, so $\mu^{\mathsf{e}}_{\mathcal{I}}(\neg U) = \mu(\varnothing) = 0$ by Definition 5.12.

Case 2: there is no A with $(A,\mathcal{I}) \in U$, so $\mu^{\mathsf{e}}_{\mathcal{I}}(U) = \mu(\varnothing) = 0$ by Definition 5.12. Then we claim $(\mathcal{I},\mathcal{I}) \in \neg U$. For if $(\mathcal{I},\mathcal{I}) \notin \neg U$, then there is some $(A',I') \between (\mathcal{I},\mathcal{I})$ with $(A',I') \in U$. From $(A',I') \between (\mathcal{I},\mathcal{I})$, we have $A' \subseteq \mathcal{I} \subseteq I'$, so $(A',\mathcal{I}) \sqsubseteq (A',I')$ by Lemma 5.2 and hence $(A',\mathcal{I}) \in U$ by Lemma 4.4, contradicting the assumption of the case. Thus, $(\mathcal{I},\mathcal{I}) \in \neg U$, so $\mu^{\mathsf{e}}_{\mathcal{I}}(\neg U) = \mu(\mathcal{I}) = 1$ by Definition 5.12.

Case 3: There is an A with $(A,\mathcal{I}) \in U$ but $(\mathcal{I},\mathcal{I}) \notin U$. Let A be the largest set of worlds such that $(A,\mathcal{I}) \in U$. We claim that $(\mathcal{I} \setminus A,\mathcal{I}) \in \neg U$ and there is no $C \subseteq \mathcal{I}$ with $\mathcal{I} \setminus A \subsetneq C$ and $(C,\mathcal{I}) \in \neg U$. Then $\mu^{\mathsf{e}}_{\mathcal{I}}(U) = \mu(A)$ and $\mu^{\mathsf{e}}_{\mathcal{I}}(\neg U) = \mu(\mathcal{I} \setminus A)$ by Definition 5.12 and Lemma 5.13. Then since $\mu(\mathcal{I}) = 1$, we have $\mu(A) = 1 - \mu(\mathcal{I} \setminus A)$, so $\mu^{\mathsf{e}}_{\mathcal{I}}(U) = 1 - \mu^{\mathsf{e}}_{\mathcal{I}}(\neg U)$. Toward a contradiction, suppose $(\mathcal{I} \setminus A,\mathcal{I}) \notin \neg U$, so for some $(A',I') \between (\mathcal{I} \setminus A,\mathcal{I})$, we have $(A',I') \in U$. We claim that $(A \cup A',\mathcal{I}) \in U$. For suppose $(A'',I'') \between (A \cup A',\mathcal{I})$. Then $A'' \cap (A \cup A') \neq \varnothing$, which implies $A'' \cap A \neq \varnothing$ or $A'' \cap A' \neq \varnothing$.

Case 3a: $A'' \cap A \neq \varnothing$. Then $(A'',I'') \between (A \cup A',\mathcal{I})$ implies $(A'',I'') \between (A,\mathcal{I})$, which with $(A,\mathcal{I}) \in U$ implies there is an $(A''',I''') \between (A'',I'')$ with $(A''',I''') \in U$.

Case 3b: $A'' \cap A = \varnothing$ but $A'' \cap A' \neq \varnothing$. Since $A'' \cap A = \varnothing$, plus $A'' \subseteq \mathcal{I}$ from $(A'',I'') \between (A \cup A',\mathcal{I})$, we have $A'' \subseteq \mathcal{I} \setminus A$. Then since $\mathcal{I} \setminus A \subseteq I'$ from $(A',I') \between (\mathcal{I} \setminus A,\mathcal{I})$, we have $A'' \subseteq I'$. Then since $A'' \cap A' \neq \varnothing$, plus $A' \subseteq I''$ from $(A'',I'') \between (A \cup A',\mathcal{I})$, we have $(A'',I'') \between (A',I')$.

Thus, we have shown that for every $(A'',I'') \between (A \cup A',\mathcal{I})$, there is an $(A''',I''') \between (A'',I'')$ with $(A''',I''') \in U$. Therefore, $(A \cup A',\mathcal{I}) \in U$. But since we have $A' \cap (\mathcal{I} \setminus A) \neq \varnothing$ from $(A',I') \between (\mathcal{I} \setminus A,\mathcal{I})$, it follows that $A \subsetneq A \cup A'$. But then $(A \cup A',\mathcal{I}) \in U$ contradicts the fact that A is the largest set of worlds such that $(A,\mathcal{I}) \in U$. Thus, we conclude that $(\mathcal{I} \setminus A,\mathcal{I}) \in \neg U$. Finally, if $C \subseteq \mathcal{I}$ with $\mathcal{I} \setminus A \subsetneq C$, then $C \cap A \neq \varnothing$, which implies $(A,\mathcal{I}) \between (C,\mathcal{I})$, so $(C,\mathcal{I}) \notin \neg U$, which completes the proof of part 2. □

The classical property that certainty is preserved by conjunction is also satisfied, provided we assign *nonzero* probability to every outcome that *might* obtain (otherwise one will have credence 1 that an outcome *will not* obtain and yet also have credence 1 that it *might*, while assigning zero credence to the conjunction of these propositions, since the conjunction is an epistemic contradiction).

Proposition 5.19. For any W, $\mathcal{I}$, and μ as in Definition 5.12, the following are equivalent:

1. $\mathcal{I} = \{w \in W \mid \mu(\{w\}) > 0\}$;

2. for all $U, V \in O(\wp(W)^{\mathsf{e}})$, if $\mu^{\mathsf{e}}_{\mathcal{I}}(U) = 1$ and $\mu^{\mathsf{e}}_{\mathcal{I}}(V) = 1$, then $\mu^{\mathsf{e}}_{\mathcal{I}}(U \wedge V) = 1$.

Proof. Suppose 1 holds. Then if $\mu^{\mathsf{e}}_{\mathcal{I}}(U) = 1$ and $\mu^{\mathsf{e}}_{\mathcal{I}}(V) = 1$, we claim $(\mathcal{I},\mathcal{I}) \in U \cap V$. For if $A \subsetneq \mathcal{I}$ for the largest A such that $(A,\mathcal{I}) \in U$, then $\mu^{\mathsf{e}}_{\mathcal{I}}(U) = \mu(A) \neq 1$ by 1, contradicting our assumption, and similarly for V. Then since $(\mathcal{I},\mathcal{I}) \in U \cap V = U \wedge V$, we have $\mu^{\mathsf{e}}_{\mathcal{I}}(U \wedge V) = \mu(\mathcal{I}) = 1$.

Now suppose 1 fails. Since $\mu(\mathcal{I}) = 1$, we have $\mathcal{I} \supsetneq \{w \in W \mid \mu(\{w\}) > 0\}$, so it follows that there is some $w \in \mathcal{I}$ such that $\mu(\{w\}) = 0$. Hence $\mu^{\mathsf{e}}_{\mathcal{I}}(\neg e(\{w\})) = \mu(W \setminus \{w\}) = 1$. Yet since $w \in \mathcal{I}$, we have $(\mathcal{I},\mathcal{I}) \in \Diamond e(\{w\})$ and hence $\mu^{\mathsf{e}}_{\mathcal{I}}(\Diamond e(\{w\})) = 1$. But $\mu^{\mathsf{e}}_{\mathcal{I}}(\neg e(\{w\}) \wedge \Diamond e(\{w\})) = \mu^{\mathsf{e}}_{\mathcal{I}}(\varnothing) = 0$, so 2 fails. □

A kind of measure on ortholattices that (a) restricts to a finitely additive probability measure on a distinguished Boolean subalgebra, as in Corollary 5.17, and (b) satisfies the properties in Propositions 5.18

and 5.19.2 for the whole ortholattice has been studied in the mathematical literature inspired by quantum mechanics. In particular, Corollary 5.17 and Proposition 5.18 show that $\mu^{\mathsf{e}}_{\mathcal{I}}$ is a *partially additive measure* in the sense of Tkadlec 1993, Def. 2.1 and Tkadlec 1991, Def. 1.2, while Proposition 5.19 shows that it is a partially additive *Jauch-Piron* measure (Tkadlec 1993, Def. 2.1) if the condition in Proposition 5.19.1 holds.

In fact, we can go beyond property (a) above and prove that the epistemic extension of a measure restricts to a finitely additive probability measure on all of the Boolean subalgebras $\mathbb{B}_n$, not just $\mathbb{B}_0$.

Proposition 5.20. Let W, $\mathcal{I}$, and μ be as in Definition 5.12 with $\mathcal{I} = \{w \in W \mid \mu(\{w\}) > 0\}$. Let $\mathbb{B}_0$ be the image of the embedding e of $\wp(W)$ into $O(\wp(W)^{\mathsf{e}})$ from Theorem 5.7.2. Then for any $\mathbb{B}_n$ in the hierarchy of Definition 4.40, $\mu^{\mathsf{e}}_{\mathcal{I}}$ restricted to $\mathbb{B}_n$ is a finitely additive probability measure.

Proof. First observe that if there is any A with $(A, \mathcal{I}) \in \Box U$, which implies $(A, \mathcal{I}) \in \Box\Box U$ by Theorem 5.7.3, then since $(A, \mathcal{I})R(\mathcal{I}, \mathcal{I})$, we have $(\mathcal{I}, \mathcal{I}) \in \Box U$. Thus, either there is no A with $(A, \mathcal{I}) \in \Box U$, in which case $\mu^{\mathsf{e}}_{\mathcal{I}}(\Box U) = 0$, or $(\mathcal{I}, \mathcal{I}) \in \Box U$, in which case $\mu^{\mathsf{e}}_{\mathcal{I}}(\Box U) = 1$. Hence the generators of $\mathbb{B}_1$, those $\Box U$ for $U \in \mathbb{B}_0$, have measure 0 or 1. It follows by Propositions 5.18 and 5.19 that all $U \in \mathbb{B}_1$ have measure 0 or 1.

Now suppose $U, V \in \mathbb{B}_1$ and $U \wedge V = \varnothing$. Then we cannot have both $\mu^{\mathsf{e}}_{\mathcal{I}}(U) = 1$ and $\mu^{\mathsf{e}}_{\mathcal{I}}(V) = 1$, for this would imply $\mu^{\mathsf{e}}_{\mathcal{I}}(U \wedge V) = 1$ by Proposition 5.19, contradicting $U \wedge V = \varnothing$. If $\mu^{\mathsf{e}}_{\mathcal{I}}(U) = 1$ and $\mu^{\mathsf{e}}_{\mathcal{I}}(V) = 0$, or $\mu^{\mathsf{e}}_{\mathcal{I}}(U) = 0$ and $\mu^{\mathsf{e}}_{\mathcal{I}}(V) = 1$, then by Proposition 5.18.1, $\mu^{\mathsf{e}}_{\mathcal{I}}(U \vee V) = 1$, so finite additivity is not violated. Finally, suppose $\mu^{\mathsf{e}}_{\mathcal{I}}(U) = 0$ and $\mu^{\mathsf{e}}_{\mathcal{I}}(V) = 0$. Then by Proposition 5.18.2, $\mu^{\mathsf{e}}_{\mathcal{I}}(\neg U) = 1$ and $\mu^{\mathsf{e}}_{\mathcal{I}}(\neg V) = 1$, so by Proposition 5.19, $\mu^{\mathsf{e}}_{\mathcal{I}}(\neg U \wedge \neg V) = 1$, so by Proposition 5.18.2 again, $\mu^{\mathsf{e}}_{\mathcal{I}}(\neg(\neg U \wedge \neg V)) = 0$ and hence $\mu^{\mathsf{e}}_{\mathcal{I}}(U \vee V) = 0$, so again finite additivity is not violated.

Since finite additivity is not violated on $\mathbb{B}_1$, and $\mathbb{B}_n = \mathbb{B}_1$ for all $n \geq 1$ as in the proof of Theorem 5.10, and finite additivity is not violated on $\mathbb{B}_0$ by Corollary 5.17, we are done. ☐

Thus, not only classical logic but also classical probability holds for propositions at a given epistemic level.

Remark 5.21. Let us say that a function μ from an epistemic ortho-Boolean lattice L to $[0, 1]$ is an *epistemic measure* if (i) for all $a, b \in L$, $a \leq b$ implies $\mu(a) \leq \mu(b)$, (ii) $\mu(\neg a) = 1 - \mu(a)$, (iii) $\mu(a) = 1$ and $\mu(b) = 1$ jointly imply $\mu(a \wedge b) = 1$, and (iv) the restriction of μ to each $\mathbb{B}_n$ is a finitely additive probability measure. Thus, we showed above that the epistemic extension $\mu^{\mathsf{e}}_{\mathcal{I}}$ is an epistemic measure (under a regularity assumption on $\mathcal{I}$). Moreover, we think the properties of epistemic measures are all plausible when we allow functions, unlike $\mu^{\mathsf{e}}_{\mathcal{I}}$, that assign probabilities other than 0 or 1 to epistemic modal propositions. We leave for future work the study of epistemic measures with non-extreme values for modal propositions.

5.3 Lifting conditionals

In this section, we show how a conditional operator on sets-of-worlds propositions can be extended to a conditional operator on propositions in the associated epistemic frame. From a general point of view, this further demonstrates the broad applicability of our epistemic extension construction. More particularly, (indicative) conditionals have a special relationship with epistemic modality, as many have noted (see especially Kratzer 1981). One way to bring this out is to note that in the scope of a conditional antecedent, the antecedent seems to be assumed not only to be true but to be epistemically necessary. For example, given that every red card in this deck of cards is either a heart or diamond, it seems that you can be sure, not only of (25-a), but also of (25-b):

(25) a. If Latif drew a red card, it is hearts or diamonds.

b. If Latif drew a red card, it must be hearts or diamonds.

This point is closely related to failures of Modus Tollens (the inference from $\{p \to q, \neg q\}$ to $\neg p$) for conditionals with modal consequents discussed in Veltman 1985 and Yalcin 2012c. If you are ignorant of how the card draw came up, then in addition to being sure of (25-b), you should be sure that it might not be hearts or diamonds: that is, you should be sure of the negation of the consequent of (25-b). Nonetheless, you cannot use modus tollens to conclude that the card is not red.

These two points look closely related to two corresponding observations about disjunction that have motivated our possibility semantics. First, $p \vee \Box\neg p$ is always true, since it is equivalent to the negation of a Wittgenstein sentence. Second, disjunctive syllogism thus fails, since the inference from $\{p \vee \Box\neg p, \Diamond p\}$ to p is invalid. For instance, just as (25-b) is intuitively true in this situation, so is 'Either the card is black or it must be hearts or diamonds'; likewise, the inference from the latter together with 'It might be black' to 'It's black' is clearly invalid. Hence it is natural to look for a parallel account of these two phenomena in the framework of possibility semantics.

As our baseline possible worlds semantics for a conditional, we will start with Stalnaker's selection function semantics (Stalnaker 1968). Other baseline semantics could be lifted to our system in a similar way, though with different results; we are sympathetic to a broadly Stalnakerian starting point but will not do much to motivate it here, as our main focus is on the lifting construction. Recall that in Stalnaker's framework, a conditional $p \to q$ is evaluated at a world w by selecting the closest possible world to w where p is true (if there is one) and checking whether q is true there. More generally, following Lewis (1973), one might select a set of closest worlds where p is true and check whether q is true at all of them. Formally, given a set W and set-selection function $f : (W \times \wp(W)) \to \wp(W)$, we define a binary operation $\to_f$ on $\wp(W)$ by

$$A \to_f B = \{w \in W \mid f(w, A) \subseteq B\}. \tag{V}$$

We will consider ways of extending f to a set-selection functions $g : (S \times O(\wp(W)^{\mathsf{e}})) \to \wp(S)$ for the epistemic frame $\wp(W)^{\mathsf{e}} = (S, \between, R)$. From such a set-selection function, we define binary operations $\to_g$ and $\Rightarrow_g$ on $O(\wp(W)^{\mathsf{e}})$ as follows:

$$\mathscr{U} \to_g \mathscr{V} \quad = \quad \{x \in S \mid g(x, \mathscr{U}) \subseteq \mathscr{V}\} \tag{VI}$$
$$\mathscr{U} \Rightarrow_g \mathscr{V} \quad = \quad \{x \in S \mid \forall x' \between x \; \exists x'' \between x' : g(x, \mathscr{U}) \subseteq \mathscr{V}\}. \tag{VII}$$

We say that g is *regular* if for all $\mathscr{U}, \mathscr{V} \in O(\wp(W)^{\mathsf{e}})$, $\mathscr{U} \to_g \mathscr{V} = \mathscr{U} \Rightarrow_g \mathscr{V}$, which is equivalent to the condition that $\mathscr{U} \to_g \mathscr{V}$ is a $\between$-regular set.[30] Even if g is not regular, we have that $\mathscr{U} \Rightarrow_g \mathscr{V}$ is $\between$-regular.

Finally, let us add an indicative conditional $\to$ to our formal language, interpreted by:

- $\mathcal{M}, x \Vdash \varphi \to \psi$ iff $x \in [\![\varphi]\!]^{\mathcal{M}} \to_g [\![\psi]\!]^{\mathcal{M}}$

in models where g is regular and otherwise with $\Rightarrow_g$ in place of $\to_g$.

5.3.1 Nonmodal antecedents

Our goal is to relate a selection function g for possibilities to an underlying selection function f for possible worlds. We start by considering the case of non-modal antecedents, that is, where g takes as its argument

[30] Equivalently, regularity is the condition that for all $\mathscr{U} \in O(\wp(W)^{\mathsf{e}})$ and $x, y \in S$, if some possibility in $g(x, \mathscr{U})$ is compatible with y, then $\exists x' \between x \; \forall x'' \between x'$ some possibility in $g(x'', \mathscr{U})$ is compatible with y.

a possibility together with a proposition in the image of a possible worlds proposition under our epistemic extension embedding.

Definition 5.22. Let W be a nonempty set, $f : (W \times \wp(W)) \to \wp(W)$ a set-selection function, and S the set of possibilities in the epistemic extension $\wp(W)^{\mathsf{e}}$. Then a set-selection function $g : (S \times O(\wp(W)^{\mathsf{e}})) \to \wp(S)$ is *an epistemic extension of* f if for all nonempty $C \subseteq W$, we have

$$g((A, I), e(C)) = \{(\bigcup\{f(w, C) \mid w \in A\}, \bigcup\{f(w, C) \mid w \in I\})\} \tag{VIII}$$

where e is the embedding from Theorem 5.7.2.

Informally, the definition states that given a possibility (A, I) and the possibilistic projection $e(C)$ of a worldly proposition C, we get a unique selected possibility (A', I') in $g((A, I), e(C))$ by (i) applying the original set-selection function f to each world in A together with C, and unioning the results to get A' and then (ii) applying f to each world in I together with C and unioning the results to get I'.

Proposition 5.23. If W, f, and g are as in Definition 5.22, then the embedding e from Theorem 5.7.2 also preserves the conditional, i.e., for all $C, D \in \wp(W)$:

$$e(C \to_f D) = e(C) \to_g e(D) = e(C) \Rightarrow_g e(D).$$

Proof. For all $C, D \in \wp(W)$, we have:

$$\begin{aligned}
e(C \to_f D) &= \{(A, I) \in S \mid A \subseteq C \to_f D\} \text{ by definition of } e\\
&= \{(A, I) \in S \mid \text{for all } w \in A,\ f(w, C) \subseteq D\} \text{ by definition of } \to_f\\
&= \{(A, I) \in S \mid \{(\bigcup\{f(w, C \mid w \in A\}, \bigcup\{f(w, C) \mid w \in I\})\} \subseteq e(D)\} \text{ by definition of } e\\
&= \{(A, I) \in S \mid g((A, I), e(C)) \subseteq e(D)\} \text{ since } g \text{ is an epistemic extension of } f\\
&= e(C) \to_g e(D) \text{ by definition of } \to_g .
\end{aligned}$$

Then since $e(C \to_f D)$ is $\between$-regular by Theorem 5.7.2, we have $e(C) \to_g e(D) = e(C) \Rightarrow_g e(D)$. ☐

A desirable prediction of this approach to modals and conditionals is the following scopelessness property, which implies that for non-modal φ and ψ, $\Box(\varphi \to \psi)$ is equivalent to $\varphi \to \Box\psi$.

Proposition 5.24. If W, f, g, and e are as in Proposition 5.23, then for all $C, D \in \wp(W)$,

$$\Box(e(C) \to_g e(D)) = e(C) \to_g \Box e(D).$$

Proof. We have

$$\begin{aligned}
\Box(e(C) \to_g e(D)) &= \Box e(C \to_f D) \text{ by Proposition 5.23}\\
&= \{(A, I) \in S \mid I \subseteq C \to_f D\} \text{ by Lemma 5.8}\\
&= \{(A, I) \in S \mid \bigcup\{f(w, C) \mid w \in I\} \subseteq D\} \text{ by definition of } \to_f\\
&= \{(A, I) \in S \mid g((A, I), e(C)) \subseteq \Box e(D)\} \text{ by (VIII) and Lemma 5.8}\\
&= e(C) \to \Box e(D) \text{ by definition of } \to_g .
\end{aligned}$$ ☐

Corollary 5.25. Let W, f, g, and e be as in Proposition 5.23, and let S be the set of possibilities of $\wp(W)^{\mathsf{e}}$. If f satisfies the *success postulate* that $f(w, C) \subseteq C$ for all $w \in W$, then for any $C \subseteq D \subseteq W$, we have:

$$e(C) \to \Box e(D) = S.$$

Proof. If f satisfies the success postulate and $C \subseteq D$, then $C \to_f D = W$, so by Proposition 5.23, we have $e(C) \to_g e(D) = e(C \to_f D) = e(W) = S$, which implies $\Box(e(C) \to_g e(D)) = \Box S = S$ and hence $e(C) \to_g \Box e(D) = S$ by Proposition 5.24. □

Proposition 5.24 shows that our lifted conditional operator, interpreted as the indicative conditional, accounts for the key interaction between indicative conditionals and epistemic modals noted above. Since it must be that if a red card was selected, then it is diamonds or hearts, it follows that if a red card was selected, it must be diamonds or hearts. And this, in turn, shows that modus tollens will fail for modal consequents: in the epistemic model coming from a possible worlds model where every red card is diamonds or hearts, 'if a red card was selected, it must be diamonds or hearts' is true at every possibility, but by part 4 of Theorem 5.7, 'It might not be diamonds or hearts' does not entail 'it's not diamonds or hearts'.

5.3.2 From sequences of worlds to possibilities

For a concrete illustration of the construction in § 5.3.1, we need to start with a Stalnaker frame (W, f). A convenient way to construct plausible Stalnaker frames is to take the "worlds" to be *sequences* of worlds from a prior set. This follows the tradition of van Fraassen 1976 and simplifies the construction by letting "worlds" already encode information about relative closeness, rather than treating this as a separate parameter.

Definition 5.26. Given a countable set W of worlds, let W^* be the set of all sequences (indexed by an initial segment of $\mathbb{N}$) that list all elements of W without repetition. Let $s_{\wp(W)}$ be the embedding of $\wp(W)$ into $\wp(W^*)$ given by $s_{\wp(W)}(A) = \{s \in W^* \mid s \text{ starts with } w\}$. Conversely, given a proposition $\mathcal{U} \subseteq W^*$, let

$$\mathcal{U}_{\downarrow} = \{w \in W \mid \text{some sequence in } \mathcal{U} \text{ starts with } w\}.$$

Define a set-selection function $f : (W^* \times \wp(W^*)) \to W^*$ as follows:

1. $f(s, \mathcal{A}) = \varnothing$ if $\mathcal{A} = \varnothing$;
2. otherwise $f(s, \mathcal{A})$ is the singleton set of the sequence obtained from s by putting all worlds in $\mathcal{A}_{\downarrow}$, ordered as in s, before all worlds not in $\mathcal{A}_{\downarrow}$, ordered as in s.[31]

Example 5.27. If $W = \{0, 1, 2\}$, $s = (2, 1, 0)$, and $\mathcal{A} = s_{\wp(W)}(\{0, 1\})$, then $f(s, \mathcal{A}) = \{(1, 0, 2)\}$.

This construction also leads to a simple lifting of probabilities from $\wp(W)$ to $\wp(W^*)$, based on Khoo and Santorio's (2018) version of van Fraassen's (1976) construction, which we can then lift to probabilities on $O(\wp(W^*)^{\mathsf{e}})$ as in § 5.2.

Definition 5.28. For finite W, given a probability measure μ on $\wp(W)$, let μ^* be the measure on $\wp(W^*)$ such that the probability of a sequence $s \in W^*$ is the probability of obtaining s by sampling without replacement from W according to μ.

[31] In the belief revision literature, this way of reordering s is known as "lexicographic revision" or "radical upgrade" with $\mathcal{A}_{\downarrow}$. This approach, while equivalent to Stalnaker's for conditionals whose antecedents and consequents are modal- and conditional-free, diverges in some cases for complex conditionals; we defer a rigorous comparison of the approaches to future work.

Now we have everything we need for an interesting example.

Example 5.29. As in Example 5.5, suppose we are about to roll a three-sided die, so $W = \{0,1,2\}$. Then

$$W^* = \{(0,1,2),(0,2,1),(1,0,2),(1,2,0),(2,0,1),(2,1,0)\}.$$

Let f be the selection function defined as in Definition 5.26. Let g be the epistemic extension of f that gives empty output for antecedents not in the image of the embedding e of $\wp(W^*)$ into $O(\wp(W^*)^{\mathsf{e}})$ from Theorem 5.7.2:

$$g((A,I),\mathscr{U}) = \begin{cases} \{\big(\bigcup\{f(w,e^{-1}(\mathscr{U})) \mid w \in A\}, \bigcup\{f(w,e^{-1}(\mathscr{U})) \mid w \in I\}\big)\} & \text{if } \mathscr{U} \in e[\wp(W^*)] \\ \varnothing & \text{otherwise.} \end{cases}$$

Finally, let μ be the uniform measure on W, so μ^* is the lifted measure on $\wp(W^*)$ as in Definition 5.28, and $(\mu^*)^{\mathsf{e}}$ is the lifted measure on $O(\wp(W^*)^{\mathsf{e}})$ as in Definition 5.12. Let $\mathtt{0},\mathtt{1},\mathtt{2}$ be propositional variables, and define a valuation $V : \{\mathtt{0},\mathtt{1},\mathtt{2}\} \to O(\wp(W^*)^{\mathsf{e}})$ by

$$V(\mathtt{n}) = e(s_{\wp(W)}(\{n\})).$$

Then we have a possibility model $\mathcal{M} = (\wp(W^*)^{\mathsf{e}}, V)$ and the following probabilities according to $(\mu^*)^{\mathsf{e}}$:

- Probability of $[\![\mathtt{0}]\!]^{\mathcal{M}}$ is $1/3$;
- Probability of $[\![(\mathtt{0} \vee \mathtt{1}) \to \mathtt{0}]\!]^{\mathcal{M}}$ is $1/2$;
- Probability of $[\![\Diamond((\mathtt{0} \vee \mathtt{1}) \to \mathtt{0})]\!]^{\mathcal{M}}$ is 1;
- Probability of $[\![\Box((\mathtt{0} \vee \mathtt{1}) \to \mathtt{0})]\!]^{\mathcal{M}}$ is 0;
- Probability of $[\![(\mathtt{0} \vee \mathtt{1}) \to \Box\mathtt{0}]\!]^{\mathcal{M}}$ is 0 (the formula is equivalent to $\Box((\mathtt{0} \vee \mathtt{1}) \to \mathtt{0})$);
- Probability of $[\![\neg(\mathtt{1} \vee \mathtt{2}) \to \Box\mathtt{0}]\!]^{\mathcal{M}}$ is 1 (the formula is true at all possibilities);
- Probability of $[\![\mathtt{0} \to ((\mathtt{0} \vee \mathtt{1}) \to \mathtt{0}))]\!]^{\mathcal{M}}$ is 1 (the formula is true at all possibilities);
- Probability of $[\![((\mathtt{0} \vee \mathtt{1} \vee \mathtt{2}) \to \mathtt{0}) \to \mathtt{0}]\!]^{\mathcal{M}}$ is 1 (the formula is true at all possibilities);
- Probability of $[\![\mathtt{0} \to \Diamond\neg\mathtt{0}]\!]^{\mathcal{M}}$ is 0 (the formula is true at no possibilities).

Note that the probability $(\mathtt{0} \vee \mathtt{1}) \to \mathtt{0}$ agrees with Stalnaker's Thesis that the probability of a conditional is the probability of the consequent conditional on the antecedent (Stalnaker 1970). This is simply a consequence of starting with the selection function induced on sequences of worlds as in Definition 5.26. It is well known that this kind of selection function yields Stalnaker's Thesis for conditionals with Boolean antecedents (van Fraassen 1976). Then as a result of the embedding given by Theorem 5.7, Proposition 5.16, and Proposition 5.23, the probability of a conditional in the sequence model is the same as its probability in the epistemic frame coming from the sequence model.

5.3.3 Modal antecedents

Definition 5.22 only tells us how to handle conditionals with non-modal antecedents. It does not tell us what should be the value of, e.g., $g((A, I), \Box\mathscr{U} \vee \Box\mathscr{V})$. Of course a modal proposition like $\Box\mathscr{U} \vee \Box\mathscr{V}$ gives us some worldly information, namely $\mathscr{U} \vee \mathscr{V}$ (assuming $\mathscr{U}, \mathscr{V}$ are in the image of the embedding of $\wp(W)$ into $O(\wp(W)^{\mathsf{e}})$). In general, given an arbitrary proposition $\mathscr{U} \in O(\wp(W)^{\mathsf{e}})$, we define its *worldly projection* as

$$\mathscr{U}_{\Downarrow} = \bigcup\{A \subseteq W \mid \exists I : (A, I) \in \mathscr{U}\}.$$

Could conditioning with a modal proposition be equivalent to conditioning with its worldly projection? Not quite, since there is no guarantee that the possibility

$$\left(\bigcup\left\{f(w, (\Box\mathscr{V}_1 \vee \Box\mathscr{V}_2)_{\Downarrow}) \mid w \in A\right\}, \bigcup\left\{f(w, (\Box\mathscr{V}_1 \vee \Box\mathscr{V}_2)_{\Downarrow}) \mid w \in I\right\}\right)$$

belongs to $\Box\mathscr{V}_1 \vee \Box\mathscr{V}_2$, as $(\Box\mathscr{V}_1 \vee \Box\mathscr{V}_2)_{\Downarrow} = (\mathscr{V}_1 \vee \mathscr{V}_2)_{\Downarrow}$, and conditioning with $(\mathscr{V}_1 \vee \mathscr{V}_2)_{\Downarrow}$ (e.g., the coin landed heads or tails) does not necessarily take us to a possibility where $\Box\mathscr{V}_1 \vee \Box\mathscr{V}_2$ is true. But in order to validate the compelling Identity principle of conditional logic, $\varphi \to \varphi$, we need to ensure that $g((A, I), \mathscr{U}) \subseteq \mathscr{U}$ for any (A, I) and $\mathscr{U}$. In fact, in the spirit of Corollary 5.25, the possibilities in $g((A, I), \mathscr{U})$ should be possibilities where not only $\mathscr{U}$ is true but also where $\mathscr{U}$ *must* be true. Of course, if $\mathscr{U}$ is a fully modalized proposition, such as $\Box\mathscr{V}_1 \vee \Box\mathscr{V}_2$, then we already have $\mathscr{U} = \Box\mathscr{U}$ due to the S5 properties from Theorem 5.7.3, but if $\mathscr{U}$ is a mixed proposition such as $\mathscr{V}_1 \wedge \Diamond\mathscr{V}_2$, then we may have $\mathscr{U} \neq \Box\mathscr{U}$, in which case it is not redundant to want $g((A, I), \mathscr{U}) \subseteq \Box\mathscr{U}$.

In light of these points, a natural approach to dealing with modal antecedents $\mathscr{U}$ is the following:

- $f((A, I), \mathscr{U})$ = the set of possibilities in $\Box\mathscr{U}$ that are "closest" to the possibility

$$\left(\bigcup\{f(w, \mathscr{U}_{\Downarrow}) \mid w \in A\}, \{\bigcup f(w, \mathscr{U}_{\Downarrow}) \mid w \in I\}\right).$$

More formally, let us treat closeness in terms of minimizing a distance function.

Definition 5.30. Given W, f, and S as in Definition 5.22 and $d : (S \times S) \to \mathbb{R}_{\geq 0}$, we define a set-selection function $f^d : (S \times O(\wp(W)^{\mathsf{e}})) \to \wp(S)$ by

$$f^d((A, I), \mathscr{U}) = \underset{(A', I') \in \Box\mathscr{U}}{\arg\min}\; d\Big((A', I'),\ \left(\bigcup\left\{f(w, \mathscr{U}_{\Downarrow}) \mid w \in A\right\}, \bigcup\{f(w, \mathscr{U}_{\Downarrow}) \mid w \in I\}\right)\Big).$$

Under weak assumptions about f and d, we have that f^d is an epistemic extension of f.

Lemma 5.31. Let W, f, and S be as in Definition 5.22. If f satisfies the success postulate (recall Corollary 5.25) and for all $x, y \in S$, $d(x, y) = 0$ iff $x = y$, then f^d is an epistemic extension of f.

Proof. To see that f^d satisfies equation (VIII), Corollary 5.25 implies that for any $\mathscr{U}$ in the image of embedding e, setting $x = \left(\bigcup\left\{f(w, \mathscr{U}_{\Downarrow}) \mid w \in A\right\}, \bigcup\left\{f(w, \mathscr{U}_{\Downarrow}) \mid w \in I\right\}\right)$, we have $x \in \Box\mathscr{U}$, in which case $f^d((A, I), \mathscr{U}) = \{x\}$, since by assumption x is the closest possibility to itself. $\Box$

We now give an example of a distance function and some sample results.

Example 5.32. Recall that the *Hamming distance* between two sets X and Y, $d_H(X, Y)$, is the cardinality of the symmetric difference $(X \setminus Y) \cup (Y \setminus X)$. We lift this to a distance between possibilities by summing pointwise Hamming distances:

$$d_H((A,I),(A',I')) = d_H(A,A') + d_H(I,I').$$

Table 3 gives examples of Hamming distances between possibilities in the epistemic frame constructed from two worlds in Example 5.3. In fact, conveniently, the Hamming distance between two possibilities in the epistemic frame in Example 5.3 is simply the number of compatibility steps to get from one possibility to the other (a convenience that does not generalize to epistemic frames constructed from more than two worlds). The upper table in Figure 4 shows the set-selection function f^{d_H} starting from the set-selection function f for $W = \{0,1\}$ given by $f(w,W) = \{w\}$ and $f(w,\{v\}) = \{v\}$. It is easy to check that f^{d_H} is regular. The lower table in Figure 4 shows the induced conditional operation $\rightarrow_{f^{d_H}} = \Rightarrow_{f^{d_H}}$.

(A,I)	(A',I')	$d_H((A,I),(A',I'))$
$(\{0\},\{0,1\})$	$(\{0,1\},\{0,1\})$	1
$(\{0\},\{0,1\})$	$(\{0\},\{0\})$	1
$(\{0\},\{0,1\})$	$(\{1\},\{0,1\})$	2
$(\{0\},\{0,1\})$	$(\{1\},\{1\})$	3

Table 3: Examples of Hamming distances between possibilities in the epistemic frame from two worlds.

	x_1	x_2	x_3	x_4	x_5
$\varnothing = P \wedge \neg P$	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$
$\{x_1\} = \Box P$	$\{x_1\}$	$\{x_1\}$	$\{x_1\}$	$\{x_1\}$	$\{x_1\}$
$\{x_3\} = \Diamond P \wedge \Diamond\neg P$	$\{x_3\}$	$\{x_3\}$	$\{x_3\}$	$\{x_3\}$	$\{x_3\}$
$\{x_5\} = \Box\neg P$	$\{x_5\}$	$\{x_5\}$	$\{x_5\}$	$\{x_5\}$	$\{x_5\}$
$\{x_1,x_2\} = P$	$\{x_1\}$	$\{x_1\}$	$\{x_1\}$	$\{x_1\}$	$\{x_1\}$
$\{x_4,x_5\} = \neg P$	$\{x_5\}$	$\{x_5\}$	$\{x_5\}$	$\{x_5\}$	$\{x_5\}$
$\{x_1,x_5\} = \Box P \vee \Box\neg P$	$\{x_1\}$	$\{x_1\}$	$\{x_1,x_5\}$	$\{x_5\}$	$\{x_5\}$
$\{x_1,x_2,x_3\} = \Diamond P$	$\{x_1\}$	$\{x_2\}$	$\{x_3\}$	$\{x_3\}$	$\{x_3\}$
$\{x_3,x_4,x_5\} = \Diamond\neg P$	$\{x_3\}$	$\{x_3\}$	$\{x_3\}$	$\{x_4\}$	$\{x_5\}$
$S = P \vee \neg P$	$\{x_1\}$	$\{x_2\}$	$\{x_3\}$	$\{x_4\}$	$\{x_5\}$

$\rightarrow_{f^{d_H}}$	$\varnothing$	$\Box P$	$\Diamond P \wedge \Diamond\neg P$	$\Box\neg P$	P	$\neg P$	$\Box P \vee \Box\neg P$	$\Diamond P$	$\Diamond\neg P$	X
$\varnothing$	X	X	X	X	X	X	X	X	X	X
$\Box P$	$\varnothing$	X	$\varnothing$	$\varnothing$	X	$\varnothing$	X	X	$\varnothing$	X
$\Diamond P \wedge \Diamond\neg P$	$\varnothing$	$\varnothing$	X	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	X	X	X
$\Box\neg P$	$\varnothing$	$\varnothing$	$\varnothing$	X	$\varnothing$	X	X	$\varnothing$	X	X
P	$\varnothing$	X	$\varnothing$	$\varnothing$	X	$\varnothing$	X	X	$\varnothing$	X
$\neg P$	$\varnothing$	$\varnothing$	$\varnothing$	X	$\varnothing$	X	X	$\varnothing$	X	X
$\Box P \vee \Box\neg P$	$\varnothing$	P	$\varnothing$	$\neg P$	P	$\neg P$	X	P	$\neg P$	X
$\Diamond P$	$\varnothing$	$\Box P$	$\Diamond\neg P$	$\varnothing$	P	$\varnothing$	$\Box P$	X	$\Diamond\neg P$	X
$\Diamond\neg P$	$\varnothing$	$\varnothing$	$\Diamond P$	$\Box\neg P$	$\varnothing$	$\neg P$	$\Box\neg P$	$\Diamond P$	X	X
S	$\varnothing$	$\Box P$	$\Diamond P \wedge \Diamond\neg P$	$\Box\neg P$	P	$\neg P$	$\Box P \vee \Box\neg P$	$\Diamond P$	$\Diamond\neg P$	X

Table 4: Above: set-selection function f^{d_H} for the epistemic frame of Example 5.3 given the set-selection function f for $W = \{0,1\}$ given by $f(w,W) = \{w\}$ and $f(w,\{v\}) = \{v\}$. Below: conditional operation determined by f^{d_H}.

If we start with W^* and f as in Example 5.29, then f^{d_H} is not regular, but this only means that we must interpret the conditional $\rightarrow$ with $\Rightarrow_{f^{d_H}}$ instead of $\rightarrow_{f^{d_H}}$. Then using the same probability measure as in Example 5.29, we obtain a number of plausible results[32]:

- Probability of $[\![(\Box\mathtt{0} \vee \Box\mathtt{1}) \to \Box\mathtt{0}]\!]^{\mathcal{M}}$ is 1/2;
- Probability of $[\![\Box(\mathtt{0} \vee \mathtt{1}) \to \Box\mathtt{0}]\!]^{\mathcal{M}}$ is 0;
- Probability of $[\![(\mathtt{1} \vee \mathtt{2})]\!]^{\mathcal{M}}$ is 2/3;
- Probability of $[\![\Diamond\neg\mathtt{0} \to (\mathtt{1} \vee \mathtt{2})]\!]^{\mathcal{M}}$ is 2/3;
- Probability of $[\![\Diamond\mathtt{0} \to \neg\mathtt{0}]\!]^{\mathcal{M}}$ is 0 (the formula is true at no possibilities);
- Probability of $[\![\Diamond\mathtt{0} \to \mathtt{0}]\!]^{\mathcal{M}}$ is 1/3;
- Probability of $[\![(\mathtt{0} \wedge \Diamond\neg\mathtt{0}) \to \bot]\!]^{\mathcal{M}}$ is 1 (the formula is true at all possibilities).

Calculating these probability by hand is somewhat tedious, but code to compute them is available in the notebook on possibility semantics cited in § 1.

We leave it to future work to more systematically study different notions of distance between epistemic possibilities for the purposes of handling conditionals with modal antecedents, as well as to further explore and motivate the theory of conditionals that results from the epistemic extension of a Stalnaker conditional. Already the results above suggest that our possibility semantics provides a promising framework in which to handle the interaction of conditionals with epistemic modality and probabilities.

6 Comparisons

In this section, we compare our approach to epistemic modals to some others in the literature. We cannot cover them all, since there are far too many. We will, however, briefly situate our approach in the existing literature, with emphasis on comparing our theories to the most similar ones in the literature.[33]

Among non-E-logics, several variants are worth mentioning. First, there are logics in which Wittgenstein sentences are contradictions but cannot always be replaced by contradictions *salva veritate*. In other words, these logics have half of the profile of E-logics. A prominent example of a logic like this comes from the domain semantics for epistemic modals (Yalcin 2007, MacFarlane 2011, Bledin 2014) together with Yalcin (2007) and Bledin's (2014) notion of informational consequence or Kolodny and MacFarlane's (2010) notion

[32] While the Hamming distance d_H produces plausible probabilities, one could consider more fine-grained distance measures. For example, if possibilities are sets of *sequences* of worlds as in Example 5.29, then we can judge the distance between two possibilities (A, I) and (A', I') in terms of how "far apart" are the sequences in A and A', and similarly for I and I'. Recall that the *Kemeny distance* $d_K(s, s')$ between two sequences $s = (s_1, \dots, s_n)$ and $s' = (s'_1, \dots, s'_n)$ is the minimal number of adjacent swaps of elements that transform s into s'. For example, the Kemeny distance between $(0, 1, 2)$ and $(1, 0, 2)$ is 1, the Kemeny distance between $(0, 1, 2)$ and $(1, 2, 0)$ is 2, and the Kemeny distance between $(0, 1, 2)$ and $(2, 1, 0)$ is 3. Now given two sets A and A' of sequences, we introduce the concept of *Hamming-Kemeny distance*:

$$d_{HK}(A, A') = \sum_{s \in A} \min(\{d_K(s, s') \mid s' \in A'\}) + \sum_{s' \in A'} \min(\{d_K(s, s') \mid s \in A\}).$$

The idea of d_{HK} is that instead of incurring a fixed penalty for each $s \in A$ such that $s \notin A'$, the penalty is the minimal Kemeny distance between s and a sequence in A'; of course this minimal distance is 0 if $s \in A'$, in agreement with the 0 penalty for the Hamming distance when $s \in A'$. The same applies to each $s' \in A'$ in the other direction. For example, the Hamming-Kemeny distance between $\{(0, 1, 2)\}$ and $\{(1, 0, 2), (1, 2, 0)\}$ is $1 + (1 + 2) = 4$. As before, we lift d_{HK} to possibilities by summing pointwise distances: $d_{HK}((A, I), (A', I')) = d_{HK}(A, A') + d_{HK}(A, A')$. This leads to some subtle differences relative to d_H but agrees with d_H on all the probabilities listed below.

[33] Among those we will not discuss are the salience-based approaches of Dorr and Hawthorne 2013 and Stojnić 2016 and the probabilistic approaches of Moss 2015 and Swanson 2016. See Mandelkern 2023 for critical discussion of the former and Mandelkern 2019 and Hawke and Steinert-Threlkeld 2021 for critical discussion of the latter.

of quasi-validity. But the fact that Wittgenstein sentences cannot always be substituted for contradictions ends up depriving approaches like this of a great deal of empirical coverage, since we cannot conclude that φ and φ' are always equivalent when φ embeds a sentence $\psi \wedge \lozenge\neg\psi$ and φ' replaces that sentence with $\psi \wedge \neg\psi$. More concretely, in the examples from § 2.1 of modalized, disjoined, and quantified Wittgenstein sentences ((5)–(8)), the cited approaches fail to predict any infelicity (see Mandelkern 2019 for extensive discussion). Another way to put this is in terms of logic. For instance, in the logics of the cited approaches, $\lozenge p \wedge \lozenge\neg p$ entails $(\lozenge p \wedge \neg p) \vee (\lozenge\neg p \wedge p)$. Likewise, as Justin Helms pointed out to us, $\lozenge p$ entails $p \vee (\neg p \wedge \lozenge p)$, which is intuitively equivalent to p (see Dorr and Hawthorne 2013 for a related observation). This brings out the importance of both halves of an E-logic—a logic in which Wittgenstein sentences are contradictions *and* can always be substituted for contradictions.

A different non-E-logic with the same profile is developed in a multilateral proof-theoretic setting in Incurvati and Schlöder 2020 and Aloni et al. 2022. We are methodologically sympathetic with that approach, insofar as it focuses on directly characterizing the logic of epistemic modals. However, the fact that those are not E-logics leads to problems like those that face the domain semantics: disjoined and modalized Wittgenstein sentences are predicted to be consistent. Incurvati and Schlöder (2020) propose to explain the incoherence of those cases by adopting a pragmatic rule that says that a disjunction can only be asserted if its disjuncts are supposable, and likewise that $\lozenge\varphi$ can only be asserted if φ is supposable. Since Wittgenstein sentences are not supposable in their logic, this suffices to account for the incoherence of asserting a disjunction of Wittgenstein sentences or a Wittgenstein sentence embedded under 'might'. (A defender of domain semantics could appeal to a similar rule.)

However, this approach fall shorts when it comes to disjoined or modalized Wittgenstein sentences that are themselves embedded under attitude predicates. There is clearly no rule that says the prejacent of an attitude predicate needs to be supposable in order for the attitude ascription to be assertable. For instance, if Sue has inconsistent beliefs about the weather, we can report them as in (26):

(26) Sue believes that it is raining and that it isn't raining. Her beliefs are inconsistent!

But now consider a disjoined Wittgenstein sentence under 'believes':

(27) Sue believes that either it's raining and might not be, or it isn't raining and might be.

Our judgment about (27) is that, like (26), it ascribes incoherent beliefs to Sue. But we cannot see how an account like the one under discussion would predict this, since (i) the prejacent of 'believes' is consistent on this account; and (ii) as (26) illustrates, there is no pragmatic rule that prevents us from ascribing unsupposable beliefs to others. Similar points apply to sentences like (28):

(28) Sue believes that it might be that it's raining and might not be raining.

In our view, these points provide further evidence for an E-logic and against a pragmatic account of these data.

Another kind of non-E-logic comes from the dynamic semantic tradition (Veltman 1996, Groenendijk et al. 1996). There are many different systems under this umbrella; for concreteness, we will focus on the semantics plus consequence relation of Groenendijk et al. 1996 for the moment and then say something about how this generalizes. In the system of Groenendijk et al. 1996, sentences with the form $\varphi \wedge \lozenge\neg\varphi$ are contradictions provided that φ is Boolean. However, when φ is non-Boolean, $\varphi \wedge \lozenge\neg\varphi$ can fail to be a contradiction. And order variants like $\Diamond\varphi \wedge \neg\varphi$ are generally not contradictions. As many have observed, however, there is no evidence for the kinds of stark order asymmetries thus predicted (see, e.g., Yalcin 2012a, Mandelkern 2023). While this approach has a ready explanation of the unassertability of Wittgenstein sentences of both orders, it has a harder time making sense of various embedding data involving the left-modal variant. For instance, the standard dynamic approach to the conditional predicts that $(\Diamond\varphi \wedge \neg\varphi) \to \psi$ is coherent and indeed equivalent to $\Diamond\varphi \wedge (\neg\varphi \to \psi)$ in contexts consistent with φ. By contrast, sentences like this sound as bad as in the reverse order, as illustrated by the contrast between (29-a) and (29-b):

(29) a. #If it might be raining but it isn't, we don't need rainjackets.
b. It might be raining, but if it isn't, we don't need rainjackets.

Finally, this system invalidates many classical inferences—including the law of non-contradiction ($\varphi \wedge \neg\varphi \vdash \bot$) and excluded middle—which are intuitively valid for epistemic language (Mandelkern 2020).

There are many other approaches to epistemic modals within a broadly dynamic framework (see, e.g., van der Does et al. 1997; Aloni 2000; Yalcin 2015; Goldstein 2019; Gillies 2020), which have respectively different logics. For instance, if we combined the semantic system of Groenendijk et al. 1996 with the *test-to-test* consequence relation instead of the update-to-test relation of Groenendijk et al. 1996, we would have a logic where all Wittgenstein sentences are contradictions, as are sentences with the form $\varphi \wedge \neg\varphi$. However, this change in perspective does not yield a system that is descriptively adequate, given the desiderata laid out at the outset, because in this system contradictions do not always *embed* as we would expect contradictions to. So, for instance, while $\varphi \wedge \neg\varphi$ is a contradiction in this system, $\Diamond(\varphi \wedge \neg\varphi)$ is not, nor is $\Diamond(\Diamond\varphi \wedge \neg\varphi)$. That is, (epistemic) contradictions remain epistemically possible. Likewise, (epistemic) contradictions can be coherently embedded in the antecedents of conditionals in this system: $(\varphi \wedge \neg\varphi) \to \psi$ and $(\Diamond\varphi \wedge \neg\varphi) \to \psi$ can both be coherent and non-trivially true and false. Indeed, $(\Diamond\varphi \wedge \neg\varphi) \to \psi$ is still predicted to be equivalent to $\Diamond\varphi \wedge (\neg\varphi \to \psi)$ in contexts consistent with φ. While there are many different pairs of dynamic semantic systems and dynamic notions of consequence, with diverse logical profiles, we do not know of any that can capture the full range of embedding data sketched in § 2. (See Mandelkern 2023 for more extensive critical discussion of various dynamic systems and more on the surrounding dialectic.)

We turn now to E-logics: that is, logics that predict that Wittgenstein sentences are contradictions and are always substitutable for contradictions (in some but not all of these systems, this holds because of a more general fact that logically equivalent sentences are always substitutable *salva veritate*).

One E-logic in the literature is the bounded theory of Mandelkern 2019. That system is given in a trivalent setting, which yields (at least) two natural logics: the *truth-preserving* logic (the set of inferences that preserve truth) and the *Strawson logic* (the set of inferences that preserve truth when the premises and conclusion are all true or false). Both logics are E.[34] But neither has exactly the properties that we have argued a logic for epistemic modals should have. In the truth-preserving logic, many classical inferences, like conjunction introduction, fail; from this perspective, the logic is too non-classical. In the Strawson logic, by contrast, all classical inferences are valid. But from this perspective, the logic is too classical: inferences like distributivity and disjunctive syllogism, which intuitively fail, come out valid. From a logical perspective, neither of these profiles is exactly the right fit for epistemic modals. The combination of the two perspectives might still suffice to make sense of the data, but the present approach is logically much more perspicuous.

By identifying a logic that is sound and complete with respect to the algebraic semantics and possibility semantics we have studied, consisting of intuitively valid principles, we have strong evidence that our

[34] In the weaker of the two systems Mandelkern considers, this holds only if we restrict attention to *diagonal consequence*.

semantics does not overgenerate in its predictions of the logical inferences that will strike speakers as valid. Since capturing felt (in)validity is a central part of the task of natural language semantics, having such a characterization is an essential part of any adequate theory of a fragment of natural language. We emphasize this point not just because it brings out a contrast with the approach of Mandelkern 2019 but because in general axiomatization is often left out of contemporary semantic theory, and we think it is important methodologically to explicitly axiomatize the set of inferences predicted by a system to strike speakers as valid (see Holliday and Icard 2018 for more discussion of the role of axiomatization in semantic theory).

A different class of (nearly) E-logics is given in the state-based frameworks of Veltman 1985, Hawke and Steinert-Threlkeld 2021, Aloni 2018, and Flocke 2020, where sentences are evaluated for truth relative to information states. Formally speaking, we can (if we want) model an information state as a set of possible worlds, which in turn is a special case of a possibility. So the indices in our semantics are closely related to those in state-based semantic frameworks, which makes these approaches a direct precedent for ours. We will discuss two of these approaches in a bit more detail. First consider Veltman's *data semantics*.[35] In data semantics, $\Diamond\varphi$ is true at an information state just in case it has some refinement where φ is true. When β is Boolean, it is persistent: if β is true in an information state, it is true in every refinement of that information state. So $\beta \wedge \Diamond\neg\beta$ can never be true (conjunction has standard truth-conditions) and is everywhere substitutable for $\bot$. However, this system is not quite an E-logic because (just as in dynamic semantics) $\varphi \wedge \Diamond\neg\varphi$ is not contradictory when φ itself is epistemic. For instance, where $\varphi = \Diamond p$, $\Diamond p$ may be true in a state but false in a refinement of that state (namely, a refinement that makes $\neg p$ true), so $\Diamond p \wedge \Diamond\neg\Diamond p$ will be consistent. And Veltman's logic in general is in some ways less classical and in other ways more classical than ours—and in both directions, we think our system is an improvement. In the first direction, data semantics invalidates excluded middle (Veltman 1985, p. 172). In this respect, it is even more extreme than dynamic semantics, since excluded middle fails even for Boolean sentences (disjunction has its standard truth-conditions, so in a state that does not settle p or $\neg p$, $p \vee \neg p$ will not be true, because neither p nor $\neg p$ is true). In fact, as Veltman points out, there are no valid Boolean sentences at all. Of course, we may be able to give some other explanation of why $p \vee \neg p$ has the feeling of a validity, but it is not immediately clear what that would be. In the second direction, Veltman validates distributivity, because he validates proof by cases with side assumptions (from $\xi \wedge \varphi \vdash \chi$ and $\xi \wedge \psi \vdash \chi$ conclude $\xi \wedge (\varphi \vee \psi) \vdash \chi$), and disjunction introduction, which together entail distributivity, as noted in § 3.1.3. Thus, assuming disjunction introduction, we have an argument against proof by cases with side assumptions. For a more direct argument, note that the following are all intuitive (and hold in our logic):

$$\Diamond p \wedge \Diamond\neg p \vdash (\Diamond p \wedge \Diamond\neg p) \wedge (p \vee \neg p),\ (\Diamond p \wedge \Diamond\neg p) \wedge p \vdash \bot,\ \text{and}\ (\Diamond p \wedge \Diamond\neg p) \wedge \neg p \vdash \bot.$$

But then by proof by cases with side assumptions and transitivity of $\vdash$, we would have $\Diamond p \wedge \Diamond\neg p \vdash \bot$.

In fact, since data semantics validates distributivity and $\beta \wedge \Diamond\neg\beta$ is always substitutable for a contradiction, it has an unwanted prediction: it rightly predicts that $(\beta \wedge \Diamond\neg\beta) \vee (\neg\beta \wedge \Diamond\beta)$ is a contradiction, but then since it validates distributivity, it predicts that $(\beta \vee \neg\beta) \wedge (\Diamond\neg\beta \wedge \Diamond\beta)$ is a contradiction. This seems implausible and indeed worse than the *lack* of a prediction that $p \vee \neg p$ is valid: while the latter lack might be removed via some other theory of felt validity, it is hard to see how this approach could be made compatible with the feeling that $(\beta \vee \neg\beta) \wedge (\Diamond\neg\beta \wedge \Diamond\beta)$ is consistent and indeed equivalent to $\Diamond\neg\beta \wedge \Diamond\beta$.

Turn next to the system of Hawke and Steinert-Threlkeld 2021 (henceforth HS-T), which is designed

[35] Gauker (2020) develops an interesting system that is, as far as our concerns go, relevantly like Veltman's.

specifically to account for the epistemic modal data above and indeed does a beautiful job of capturing them—in particular, invalidating distributivity, unlike Veltman. HS-T's system, like ours, is classical over the Boolean fragment, but over the full language, it invalidates De Morgan's laws as well as disjunction introduction. Of course, we might have found failures of these laws, and HS-T indeed argue, on the basis of the free choice inference, that disjunction introduction should fail. We will not try to adjudicate that complicated topic here.[36] We think the failures of De Morgan's laws are more obviously problematic. For instance, $\Diamond p \wedge \Diamond \neg p$ is not equivalent in their system to $\neg(\Box\neg p \vee \Box p)$: while the first is true, as one would expect, in any state containing both p-worlds and $\neg p$-worlds, the second is not. In fact, for Boolean p, $\neg(\Box\neg p \vee \Box p)$ is never true. For $\neg(\Box\neg p \vee \Box p)$ is true at a state iff $\Box\neg p \vee \Box p$ is false at a state. But $\Box\neg p \vee \Box p$ is true at every state: for it is true at a state iff that state has a (possibly empty) part where $\Box\neg p$ is true and a (possibly empty) part where $\Box p$ is true. But that will always hold: we can find those parts simply by partitioning the state into the p- and $\neg p$-worlds. This prediction seems wrong. On the one hand, the sentences in (30) seem equivalent (as De Morgan's laws, and our account, predict, but contrary to HS-T); and on the other, the disjunction in (31) does not sound like a logical truth, contrary to HS-T (it feels more committal than 'Either it isn't raining or it is').

(30) a. It might be raining and it might not be raining.
b. It's not the case that it must not be raining or must be raining.

(31) It must not be raining or it must be raining.

So we think it is a substantial advantage of our view that we validate the De Morgan equivalences and do not validate $\Box\neg p \vee \Box p$.

A second point concerns probabilities. As we discussed in § 5.2, we naturally form judgments about the probabilities of sentences containing modals, which leads to interesting puzzles. As we saw, the epistemic extension of a probability measure yields a natural model of those judgments. By contrast, state-based theories face fundamental problems here. The first is conceptual. In state-based theories, at least as understood by HS-T, the indices relative to which such sentences are evaluated as true or false are information states. Those information states represent mental states, and it is unclear how we could define reasonable measures over sets of such states. (Sloganistically, the probability of a sentence is the probability that it is *true*, not the probability that it is *accepted* by some relevant agents; for a similar point, compare Khoo (2022, p. 5): 'What would it mean to assign a probability to a property of your cognitive state?') The second, closely related problem arises from the fact that in state-based theories p entails $\Box p$. This holds not just in the informational logic, but also in the logic of truth-preservation, since truth is defined relative to mental states. Reflecting just on what it takes for a mental state to *accept* something—the primitive notion in state-based semantics—this equivalence might look defensible. But as soon as we consider probabilities, it does not; the probability of $\Box p$ can be less than the probability of p. For instance, given that a fair die has been thrown, it seems reasonable to think that there is probability .5 that it landed even, but very *little* probability that it *must* have landed even. While of course it is open to state-based approaches to deny that probabilities must track entailments, it is hard to see, practically speaking, how to construct a probability measure over their

[36] HS-T also argue for the principle they call *Inheritance*, which says that $p \wedge \Diamond q \vDash \Diamond(p \wedge q)$. We showed in Proposition 5.11.3 that this principle is valid (for p, q Boolean) in epistemic frames of Boolean algebras. On the other hand, if we drop the S5 assumption built into the epistemic frame construction, then one may have doubts. For example, perhaps you can have high credence that Mary will come to the party and that Aliyah might come to the party, while having very low credence that it might be that they both come (e.g., they just broke up). Certainly Inheritance cannot be accepted for arbitrary non-Boolean p, q, since $\Diamond p \wedge \Diamond \neg p$ should not entail $\Diamond(\Diamond p \wedge \neg p)$ (i.e., $\Diamond\bot$ and hence $\bot$). Compare these points to Example 3.39.

points that avoids this conclusion, since, again, $\Box p$ is true at any state where p is. Note that this second problem will persist even if we drop HS-T's conceptualization of states as representing mental states, since it arises simply from the logic of the system. The state-based theory thus faces a fundamental problem in meshing with a reasonable account of credences defined on epistemic modal propositions.

It is worth noting that similar points apply to a range of theories that characterize sentential contents in non-standard terms, including dynamic semantics and various expressivist and relativist theories. In those systems, sentential contents in general are not the kinds of things ordinarily measured by probabilities, raising at least a prima facie worry about how such a system can make sense of exchanges like (32-a):

(32) a. The coin must have landed heads.
b. That has a .5 probability of being true.

Exchanges like this intuitively make sense (for instance, in a case where there is a .5 chance that the coin was double-headed). But in systems where the content of a sentence like (32-a) is a kind of thing that probability measures are not defined on, exchanges like this, on the face of it, will be nonsense. This is particularly problematic since many of these theories predict that while an exchange like (32-a) should make no sense, the corresponding sentence in (33) is perfectly coherent:

(33) There is .5 probability that the coin must have landed heads.

To give one example of a recent system with a profile like this, in the probabilistic semantics of Moss 2015, sentential contents are sets of probability measures, which are not, of course, themselves defined on sets of probability measures. See Khoo 2022, Goldstein 2021, and Charlow 2019 for discussion of this problem and some attempted solutions.

A final point is slightly more complicated, but it brings out an important difference between our system and the state-based ones cited above, as well as standard implementations of domain and dynamic semantics. HS-T, again, cast their system as an attempt to cash out the core of the *expressivist* approach of Yalcin 2007: namely, that to assert $\Diamond p$ is to express a certain state of mind, namely, one that is compatible with p, and thus to believe $\Diamond p$ is simply for p to be compatible with one's beliefs—hence their definition of truth at a state, understood as a mental state. This idea is tempting, but while this approach seems reasonable enough for non-factive mental states, it does not work for factive mental states (Yalcin 2012b, Dorr and Hawthorne 2013). It is not plausible that to know $\Diamond p$ is for p to be compatible with your knowledge, since any truth is compatible with your knowledge, but it does not follow from p being true that you know might p. That is, the natural extension of this approach to factives will validate the inference $p \vDash K_a \Diamond p$, but that inference is plainly invalid:

(34) The treasure is buried in Alaska, but Blackbeard doesn't even know it might be there—he's sure it's at sea.

While HS-T do not give a semantics for 'knows', they will have to contend with this problem (as will any approach built on Yalcin's central idea). While there are various responses available that complicate the semantics of attitude predicates to deal with these cases (e.g., in Yalcin 2012b, Rothschild 2011, Willer 2013, Beddor and Goldstein 2021), we by contrast will simply not face this problem in the first place. Given an accessibility relation R_{K_a} between possibilities, let K_a be the box operation defined using R_{K_a}, just as our $\Box$ is defined using R. Then if we consider, e.g., the Epistemic Scale from Example 4.33 and set R_{K_a} to be the universal relation between possibilities, representing full ignorance for agent a, then $p \wedge \neg K_a \Diamond p$ will be true at the possibilities x_1 and x_2 that make p true, so $p \not\vDash K_a \Diamond p$.

A closely connected point concerns the relation between knowledge and belief. If, as on the expressivist way of thinking, $K_a \Diamond p$ is true whenever p is compatible with a's knowledge state, and $B_a \Diamond p$ is true whenever p is compatible with a's beliefs, then $K_a \Diamond p$ will not entail $B_a \Diamond p$. Instead, $B_a \Diamond p$ will entail $K_a \Diamond p$. For whenever p is compatible with a belief state, it is compatible with the corresponding knowledge state, but not vice versa. By contrast, we can capture the correct entailment relation by simply requiring that $R_{B_a}(x) \subseteq R_{K_a}(x)$, where R_{B_a} is the doxastic accessibility relation for agent a, as in standard doxastic-epistemic logic. Then $K_a \Diamond p$ entails $B_a \Diamond p$ but not vice versa. Together with the points about logic and probability above, we take the possibility of a smooth extension of our system to factive attitude verbs to speak in favor of it over HS-T.

Finally, Kratzer (2020) develops a theory of epistemic language in a situation-based framework. Her notion of situation is very different from our notion of a possibility (see Kratzer 1989), but it is an important precedent for our approach insofar as partiality plays a central role in that theory, as in ours. Having said that, the logic of the theory is classical, with apparent failures of distributivity to be explained pragmatically, so the overall approach is very different.

7 Conclusion

Epistemic modals have strange properties. While this has been well appreciated in the recent literature, that literature has tended to take the corresponding logical failures seriously but not literally. Our goal has been to develop a logic and corresponding semantics that directly capture the apparent logical peculiarity of epistemic modals, while retaining as much of classical logic as we can. One central feature of the resulting system is that it is E: Wittgenstein sentences are contradictions and can always be substituted for contradictions *salva veritate*. That enables our system to account for the wide range of embedding data for epistemic modals, as well as the plausibly closely related data for indicative conditionals. Another feature of our system is that it invalidates distributivity, whose failure is (to a degree that has not been sufficiently appreciated) central to characterizing the logic of epistemic modals. Not only does distributivity intuitively fail, but invalidating distributivity is also equivalent (relative to a background orthologic) to rejecting the interpretation of negation as pseudocomplementation. That, in turn, allows us to block the inference from $p \wedge \Diamond \neg p \vDash \bot$ to $\Diamond \neg p \vDash \neg p$, an inference that is valid in a classical setting and that obviously would prevent us from making $p \wedge \Diamond \neg p$ a contradiction. A set of related principles, like disjunctive syllogism and orthomodularity, are also invalid in our system.

Apart from this, however, our system retains much of classical logic. The non-modal fragment is fully classical, as is each fragment consisting of sentences at the same "epistemic level"; and a wide variety of classical laws, which we have highlighted throughout, remain valid for the entire language. Our goal is to find a minimal variant on classical logic that invalidates just those laws that should fail for epistemic modals. Having said that, we have, again, no proof that we have succeeded; for there may be arguments for logics intermediate in strength between the one we have developed and classical logic. For instance, the principle we identified in Proposition 5.11 that is valid in the logic of epistemic extensions, but not in EO^+, seems plausible to us, and other weakenings of distributivity might be found that we should add to EO^+. We hope that this paper will spur exploration of such logics.

One way to taxonomize the literature on epistemic modality is in terms of where different proposals locate the central revision required by epistemic modals. The fundamental player in the dynamic treatment of Wittgenstein sentences is its non-classical treatment of conjunction (and, correspondingly, the quantifiers). The domain treatment of Wittgenstein sentences instead has a classical approach to the connectives and holds that what is responsible for the badness of Wittgenstein sentences is fundamentally intensional: the impossibility of an information state *accepting* such a sentence. The state-based approaches start from a similar intuition. The bounded theory, like the dynamic theory, locates the action again in the connectives—in particular in the treatment of conjunction and quantifiers. Our approach, by contrast, highlights a heretofore neglected option: namely, focusing on negation. That is, while we agree with the dynamic and bounded theories that the central action is in the interaction of connectives with epistemic modals, in our theory, negation plays a starring role. We think this is well motivated, for recall that we can formulate the fundamental problem without even using conjunction: we want $\{\Diamond\neg p, p\}$ to be inconsistent, but we do not want to be able to conclude $\Diamond\neg p \vDash \neg p$, as classical logic would allow us to do. By treating negation as orthocomplementation rather than pseudocomplementation, we block this inference. In fact, conjunction remains "classical" in our theory in the sense that it is just set intersection (semantically speaking) and greatest lower bound (algebraically speaking). Of course, conjunction has distinctively non-classical properties in its interaction with disjunction—most prominently, the failure of distributivity. But this is also related to our treatment of negation, since disjunction and conjunction are related by negation via the De Morgan equivalences.

Our treatment of epistemic modals also brings out the utility of working with possibilities rather than possible worlds as the building blocks of semantics. A central commitment of possible world semantics is that compatibility just is compossibility: for φ and ψ to be compatible just is for them to be true somewhere together. While this is also a commitment of possibility semantics for classical logic (recall Remark 4.9), once we take the step from possible worlds to partial possibilities, we can take the further step of pulling apart compatibility from compossibility. Indeed, pulling apart compatibility from compossibility is *the* central difference between possibility semantics for orthologic vs. classical logic. And on reflection, it is not clear that compatibility entails compossibility. Romeo and Juliet were compatible but not compossible; p and $\Diamond\neg p$ are the same. This is why they are contradictory (they are not compossible), but neither entails the negation of the other (they are compatible). These are positions we can hold only if compatibility does not entail compossibility—that is, only in the possibility framework.

We have left many open questions, and there are many points for further development. As noted, we need to axiomatize the logic of epistemic extensions of Boolean algebras. We need to develop a model of credences in epistemic modal propositions that generalizes the notion of an epistemic extension of a probability measure, in particular relaxing the S5 assumptions built into that model, which can plausibly fail in natural language. We need to further study the behavior of conditionals in possibility semantics, which includes axiomatizing conditional logics resulting from lifting selection functions from worlds to possibilities, as well as comparing our approach in detail to existing theories. We also need to extend our semantics to a broader fragment of language, most obviously to cover binary quantifiers, where (as noted briefly in § 2.2) many interesting issues arise from the interaction with epistemic modals.

In addition to all these points, a substantial further question to explore is what pragmatic system naturally goes with the possibility semantics we have developed, and, correspondingly, what the companion logic of rational acceptance ("informational logic") comes to. One natural pragmatic model builds on the treatment of conditionals we have begun to sketch here. Associate each context with a context set, as in Stalnaker 1974, but treated now as a set of possibilities rather than worlds: namely, the set of possibilities where everything commonly accepted in the conversation is true. Unlike Stalnaker, we cannot model updates intersectively, since conversants may learn $\Diamond p$ and subsequently learn $\neg p$, without this presenting any fundamental informational conflict or tension. However, the apparatus required to model conditionals yields an elegant alternate model of updates. Take the contextual selection function g, which we assume is a regular epistemic extension of a set-selection function (§ 5.3.3). Then model conversational updates (and, more generally, the evolution of attitudes) as transitions from the context set to its image under g, relative to the newly accepted content. That is, if A is asserted and accepted in a conversation whose context set is C and whose selection function is g, the updated context set is $C + A = \bigcup\{g(x, A) \mid x \in C\}$. This yields a natural corresponding informational logic: say that φ informationally entails ψ iff for any context set C, we have $x \in [\![\psi]\!]$ for all $x \in C + [\![\varphi]\!]$. Interestingly, since $p \to \Box p$ is a theorem (given Definition 5.30), this approach has an attractive feature: it predicts that φ informationally entails $\Box\varphi$, since any context updated with φ by the procedure just defined will entail $\Box\varphi$. This is in line with the informational logic proposed by Yalcin 2007 and following, where the inference from φ to $\Box\varphi$ is taken to be a central desideratum of an informational logic.

An different approach to pragmatics would be to treat the context set as in the first instance a set of worlds, which is updated in a standard way, and treat the corresponding possibility model as being derived from the context set using the method of epistemic extensions. Of course, this approach would need to be further developed to explain how we update with purely modal information.

Whether our system, or further developments of it, ultimately provides the correct account, we hope that future work will take on board our goal of directly characterizing the logical peculiarities of epistemic language, aiming to capture exactly where classical logic goes wrong and where it doesn't.

Acknowledgements

For helpful feedback, we thank Maria Aloni, Alexandru Baltag, Justin Bledin, Scott Blomgren, Kyle Blumberg, Sam Carter, Lucas Champollion, Ahmee Christensen, Yifeng Ding, J. Dmitri Gallow, Justin Helms, Alexander Kocurek, John MacFarlane, Guillaume Massas, Gareth Peter Gordon Norman, Eric Pacuit, Frank Veltman, Seth Yalcin, and participants at the LIRa seminar at the University of Amsterdam in February 2022, the NASSLLI 2022 course on Possibility Semantics in June 2022, Advances in Modal Logic 2022, the 2023 MITing of the Minds conference, the seminar on Sources of Nonclassical Logic at UC Berkeley in Spring 2023, and the Tsinghua Logic Salon in May 2023. This paper extends a proceedings paper (Holliday and Mandelkern 2022) presented at the 2022 Amsterdam Colloquium, and we thank our audience there and our two reviewers.

A Appendix

A.1 Characterization of epistemic extensions of Boolean algebras

In § 5.1, we proved that the epistemic extension of a Boolean algebra is a complete epistemic ortho-Boolean lattice. Now we characterize exactly which complete epistemic ortho-Boolean lattices arise in this way.

Definition A.1. A complete epistemic ortho-Boolean lattice L with distinguished Boolean algebra B is *extensive* if the following conditions hold:

1. L is an S5 modal ortholattice;

2. for all $x, y \in L$, if $x \not\leq y$, then there are $a, i \in B$ with $0 \neq a \leq i$ such that

$$a \wedge \bigwedge_{b \in B,\, 0 \neq b \leq a} \Diamond b \wedge \Box i \leq x \quad \text{but} \quad a \wedge \bigwedge_{b \in B,\, 0 \neq b \leq a} \Diamond b \wedge \Box i \not\leq y;$$

3. for all $a, i, a', i' \in B$ with $0 \neq a \leq i$ and $0 \neq a' \leq i$, the following are equivalent:
   (a) $a \wedge a' \neq 0$, $a \leq i'$, $a' \leq i$;
   (b) $a' \wedge \bigwedge_{b \in B,\, 0 \neq b \leq a'} \Diamond b \wedge \Box i' \not\leq \neg(a \wedge \bigwedge_{b \in B,\, 0 \neq b \leq a} \Diamond b \wedge \Box i)$;
4. for all $a, i, a', i' \in B$ with $0 \neq a \leq i$ and $0 \neq a' \leq i$, if $a' \wedge \bigwedge_{b \in B,\, 0 \neq b \leq a'} \Diamond b \wedge \Box i' \leq \bigwedge_{b \in B,\, 0 \neq b \leq a} \Diamond b \wedge \Box i$, then $a \leq a'$ and $i' \leq i$;
5. for all $a, b_k, i \in B$, if $b_k \leq a$, then $\Diamond b_1 \wedge \cdots \wedge \Diamond b_n \wedge \Box i \leq \Diamond(a \wedge \Diamond b_1 \wedge \cdots \wedge \Diamond b_n \wedge \Box i)$.

Our first task is to show that epistemic extensions of Boolean algebras are indeed extensive.

Proposition A.2. For any Boolean algebra B, its epistemic extension $O(B^{\mathsf{e}})$ is extensive.

Proof. Verifying all the parts of Definition A.1 is straightforward with the exception of part 5. Let us show that for all $U, V, W \in \mathbb{B}_0$, if $V_k \subseteq U$, then

$$\Diamond V_1 \wedge \cdots \wedge \Diamond V_n \wedge \Box W \subseteq \Diamond(U \wedge \Diamond V_1 \wedge \cdots \wedge \Diamond V_n \wedge \Box W). \tag{IX}$$

Suppose (a, i) is a possibility, so $a \leq i$, which belongs to the left-hand side. Then where e is the embedding of B into $O(B^{\mathsf{e}})$ from Theorem 5.7, we have $a \wedge e^{-1}(V_k) \neq 0$ for $1 \leq k \leq n$ and $i \leq e^{-1}(W)$ by Lemma 5.8. Now consider any $(a', i') \between (a, i)$, so $a' \leq i'$, $a \wedge a' \neq 0$, $a \leq i'$, and $a' \leq i$. Then where

$$x = \bigvee_{1 \leq i \leq n} (a \wedge e^{-1}(V_k)),$$

we have $x \leq a \leq i'$, so $a' \vee x \leq i'$. From $a' \leq i \leq e^{-1}(W)$, we have $a' \leq e^{-1}(W)$, and from $a \wedge e^{-1}(V_k) \leq i \leq e^{-1}(W)$, we have $x \leq e^{-1}(W)$, so $a' \vee x \leq e^{-1}(W)$. It follows that $(a', i') R (a' \vee x, i') \between (x, e^{-1}(W))$ and

$$(x, e^{-1}(W)) \in U \wedge \Diamond V_1 \wedge \cdots \wedge \Diamond V_n \wedge \Box W$$

by Lemma 5.8, which completes the proof that (a, i) is in the right-hand side of (IX). $\Box$

Conversely, we will show that any extensive L is isomorphic to the epistemic extension of its distinguished Boolean algebra. We begin with a useful lemma.

Lemma A.3. In any complete L as in Definition A.1, we have the following for all $a, i \in B$ with $0 \neq a \leq i$ and $x \in L$:

1. if $a \wedge \bigwedge_{b \in B,\, 0 \neq b \leq a} \Diamond b \wedge \Box i \leq \Box x$, then $\bigwedge_{b \in B,\, 0 \neq b \leq a} \Diamond b \wedge \Box i \leq \Box x$;
2. if $\bigwedge_{b \in B,\, 0 \neq b \leq a} \Diamond b \wedge \Box i \leq x$, then $\bigwedge_{b \in B,\, 0 \neq b \leq a} \Diamond b \wedge \Box i \leq \Box x$.

Proof. For part 1, we have

$$a \wedge \bigwedge_{b\in B,\, 0\neq b\leq a} \Diamond b \wedge \Box i \leq \Box x \;\Rightarrow\; \Diamond(a \wedge \bigwedge_{b\in B,\, 0\neq b\leq a} \Diamond b \wedge \Box i) \leq \Diamond\Box x \quad \text{by monotonicity of } \Diamond$$

$$\Rightarrow \bigwedge_{b\in B,\, 0\neq b\leq a} \Diamond b \wedge \Box i \leq \Diamond\Box x \leq \Box x \quad \text{by Definition A.1.5 and S5.}$$

For part 2, we have

$$\bigwedge_{b\in B,\, 0\neq b\leq a} \Diamond b \wedge \Box i \leq x \;\Rightarrow\; \Box(\bigwedge_{b\in B,\, 0\neq b\leq a} \Diamond b \wedge \Box i) \leq \Box x \quad \text{by monotonicity of } \Box$$

$$\Rightarrow \bigwedge_{b\in K} \Box\Diamond b \wedge \Box\Box i \leq \Box x \quad \text{by multiplicativity of } \Box$$

$$\Rightarrow \bigwedge_{b\in B,\, 0\neq b\leq a} \Diamond b \wedge \Box i \leq \Box x \quad \text{by S5.} \quad \Box$$

Now we prove the promised result.

Theorem A.4. If L is extensive with distinguished Boolean algebra B, then L is isomorphic to $O(B^{\mathsf{e}})$.

Proof. Let $B^{\mathsf{e}} = (S, \between, R)$. Define a map $f : L \to O(B^{\mathsf{e}})$ by

$$f(x) = \{(a, i) \in S \mid a \wedge \bigwedge_{b\in B,\, 0\neq b\leq a} \Diamond b \wedge \Box i \leq x\}.$$

That $x \not\leq y$ implies $f(x) \not\subseteq f(y)$ follows from Definition A.1.2, so f is injective.

We begin by showing that f preserves $\neg$. First, we show that $f(\neg x) \subseteq \neg f(x)$. Suppose $(a, i) \in f(\neg x)$, so

$$a \wedge \bigwedge_{b\in B,\, 0\neq b\leq a} \Diamond b \wedge \Box i \leq \neg x$$

and hence

$$x \leq \neg(a \wedge \bigwedge_{b\in B,\, 0\neq b\leq a} \Diamond b \wedge \Box i). \tag{X}$$

Now suppose $(a, i) \between (a', i')$, so $a \wedge a' \neq 0$, $a \leq i'$, and $a' \leq i$. Then by the (a) to (b) direction of Definition A.1.3, we have

$$a' \wedge \bigwedge_{b\in B,\, 0\neq b\leq a'} \Diamond b \wedge \Box i' \not\leq \neg(a \wedge \bigwedge_{b\in B,\, 0\neq b\leq a} \Diamond b \wedge \Box i),$$

which with (X) implies

$$a' \wedge \bigwedge_{b\in B,\, 0\neq b\leq a'} \Diamond b \wedge \Box i' \not\leq x,$$

so $(a', i') \notin f(x)$. It follows that $(a, i) \in \neg f(x)$.

Next we show that $\neg f(x) \subseteq f(\neg x)$. Suppose $(a, i) \notin f(\neg x)$, so

$$a \wedge \bigwedge_{b\in B,\, 0\neq b\leq a} \Diamond b \wedge \Box i \not\leq \neg x,$$

which implies

$$x \not\leq \neg(a \wedge \bigwedge_{b\in B,\, 0\neq b\leq a} \Diamond b \wedge \Box i).$$

Then by part 2 of Definition A.1, there are $a', i' \in B$ such that $0 \neq a' \leq i'$ and

$$a' \wedge \bigwedge_{b\in B,\, 0\neq b\leq a'} \Diamond b \wedge \Box i' \leq x$$

but

$$a' \wedge \bigwedge_{b\in B,\, 0\neq b\leq a'} \Diamond b \wedge \Box i' \not\leq \neg(a \wedge \bigwedge_{b\in B,\, 0\neq b\leq a} \Diamond b \wedge \Box i).$$

Then by the (b) to (a) direction of Definition A.1.3, we have $(a', i') \between (a, i)$, which with $(a', i') \in f(x)$ implies $(a, i) \notin \neg f(x)$.

Using the fact that f preserves $\neg$, we claim that $f(x)$ is a $\between$-regular set. Suppose $(a, i) \notin f(x)$, so

$$a \wedge \bigwedge_{b\in B,\, 0\neq b\leq a} \Diamond b \wedge \Box i \not\leq x$$

and hence

$$\neg x \not\leq \neg(a \wedge \bigwedge_{b\in B,\, 0\neq b\leq a} \Diamond b \wedge \Box i).$$

Then by part 2 of Definition A.1, there are $a', i' \in B$ with $0 \neq a' \leq i'$ such that

$$a' \wedge \bigwedge_{b\in B,\, 0\neq b\leq a'} \Diamond b \wedge \Box i' \leq \neg x$$

but

$$a' \wedge \bigwedge_{b\in B,\, 0\neq b\leq a'} \Diamond b \wedge \Box i' \not\leq \neg(a \wedge \bigwedge_{b\in B,\, 0\neq b\leq a} \Diamond b \wedge \Box i).$$

Then $(a', i') \in f(\neg x) = \neg f(x)$, and by Definition A.1.3, $(a', i') \between (a, i)$. This shows that $f(x)$ is $\between$-regular.

Next we observe that $f(\bigwedge_{x\in X} x) = \bigwedge_{x\in X} f(x)$:

$$\begin{aligned}
(a, i) \in f(\bigwedge_{x\in X} x) \quad &\Leftrightarrow \quad a \wedge \bigwedge_{b\in B,\, 0\neq b\leq a} \Diamond b \wedge \Box i \leq \bigwedge_{x\in X} x \\
&\Leftrightarrow \quad a \wedge \bigwedge_{b\in B,\, 0\neq b\leq a} \Diamond b \wedge \Box i \leq x \text{ for all } x \in X \\
&\Leftrightarrow \quad (a, i) \in f(x) \text{ for all } x \in X \\
&\Leftrightarrow \quad (a, i) \in \bigwedge_{x\in X} f(x).
\end{aligned}$$

Next we show that $f(\bigvee_{x\in X} x) \leq \bigvee_{x\in X} f(x)$. The converse inclusion follows from order preservation, which follows from meet preservation. Suppose $(a, i) \in f(\bigvee_{x\in X} x)$, so

$$a \wedge \bigwedge_{b\in B,\, 0\neq b\leq a} \Diamond b \wedge \Box i \leq \bigvee_{x\in X} x,$$

which implies

$$\neg(\bigvee_{x\in X} x) \leq \neg(a \wedge \bigwedge_{b\in B,\, 0\neq b\leq a} \Diamond b \wedge \Box i).$$

Further suppose that $(a', i') \between (a, i)$, so $a \wedge a' \neq 0$, $a \leq i'$, and $a' \leq i$. Then by Definition A.1.3, we have

$$a' \wedge \bigwedge_{b\in B,\, 0\neq b\leq a'} \Diamond b \wedge \Box i' \nleq \neg(a \wedge \bigwedge_{b\in B,\, 0\neq b\leq a} \Diamond b \wedge \Box i)$$

and hence

$$a' \wedge \bigwedge_{b\in B,\, 0\neq b\leq a'} \Diamond b \wedge \Box i' \nleq \neg(\bigvee_{x\in X} x) = \bigwedge_{x\in X} \neg x,$$

which implies that for some $x \in X$,

$$a' \wedge \bigwedge_{b\in B,\, 0\neq b\leq a'} \Diamond b \wedge \Box i' \nleq \neg x.$$

Thus, $(a', i') \notin f(\neg x) = \neg f(x)$, so there is an $(a'', i'') \between (a', i')$ such that $(a'', i'') \in f(x)$. This shows that $(a, i) \in \bigvee_{x\in X} f(x)$.

Next we show that f is surjective. For $U \in O(B^{\mathrm{e}})$, where

$$x_U = \bigvee_{(a,i)\in U} \big(a \wedge \bigwedge_{b\in B,\, 0\neq b\leq a} \Diamond b \wedge \Box i\big),$$

we claim that $f(x_U) = U$. By join preservation, we have

$$f(x_U) = \bigvee_{(a,i)\in U} f\big(a \wedge \bigwedge_{b\in B,\, 0\neq b\leq a} \Diamond b \wedge \Box i\big).$$

Obviously $U \subseteq f(x_U)$, since $(a, i) \in f\big(a \wedge \bigwedge_{b\in B,\, 0\neq b\leq a} \Diamond b \wedge \Box i\big)$ by definition of f. Conversely, suppose $(a^*, i^*) \in f(x_U)$ and $(a', i') \between (a^*, i^*)$. Then since $(a^*, i^*) \in f(x_U)$, there is an $(a'', i'') \between (a', i')$ such that for some $(a, i) \in U$,

$$(a'', i'') \in f\big(a \wedge \bigwedge_{b\in B,\, 0\neq b\leq a} \Diamond b \wedge \Box i\big),$$

which means

$$a'' \wedge \bigwedge_{b\in B,\, 0\neq b\leq a''} \Diamond b \wedge \Box i'' \leq a \wedge \bigwedge_{b\in B,\, 0\neq b\leq a} \Diamond b \wedge \Box i$$

and therefore

$$\neg(a \wedge \bigwedge_{b\in B,\, 0\neq b\leq a} \Diamond b \wedge \Box i) \leq \neg(a'' \wedge \bigwedge_{b\in B,\, 0\neq b\leq a''} \Diamond b \wedge \Box i'').$$

Hence any element not under the right side is not under the left, which with Definition A.1.3 implies that any possibility compatible with (a'', i'') is also compatible with (a, i), so $(a, i) \in U$ implies $(a'', i'') \in U$ by Lemma 4.4. Thus, for any $(a', i') \between (a^*, i^*)$, there is an $(a'', i'') \between (a', i')$ with $(a'', i'') \in U$, which with $U \in O(B^{\mathrm{e}})$ implies $(a^*, i^*) \in U$, which completes the proof that $f(x_U) \subseteq U$.

Next we show that $f(\Box x) \subseteq \Box f(x)$. Suppose $(a, i) \in f(\Box x)$, so

$$a \wedge \bigwedge_{b\in B,\, 0\neq b\leq a} \Diamond b \wedge \Box i \leq \Box x.$$

Then by Lemma A.3.1, we have

$$\bigwedge_{b\in B,\, 0\neq b\leq a} \Diamond b \wedge \Box i \leq \Box x.$$

Now if $(a, i)R(a', i')$, so $a \leq a'$ and $i' \leq i$, then a is among the nonzero elements of B below a', and $\Box i' \leq \Box i$, so we have

$$a' \wedge \bigwedge_{b\in B,\, 0\neq b\leq a'} \Diamond b \wedge \Box i' \leq \bigwedge_{b\in B,\, 0\neq b\leq a'} \Diamond b \wedge \Box i \leq \Box x \leq x,$$

so $(a', i') \in f(x)$. This shows that $(a, i) \in \Box f(x)$.

Finally, we show that $\Box f(x) \subseteq f(\Box x)$. Suppose $(a, i) \notin f(\Box x)$, so

$$a \wedge \bigwedge_{b\in B,\, 0\neq b\leq a} \Diamond b \wedge \Box i \nleq \Box x.$$

which by Lemma A.3.2 implies

$$\bigwedge_{b\in B,\, 0\neq b\leq a} \Diamond b \wedge \Box i \nleq x.$$

Now by part 2 of Definition A.1, there are $a', i' \in B$ such that $0 \neq a' \leq i'$ and

$$a' \wedge \bigwedge_{b\in B,\, 0\neq b\leq a'} \Diamond b \wedge \Box i' \leq \bigwedge_{b\in B,\, 0\neq b\leq a} \Diamond b \wedge \Box i \tag{XI}$$

but

$$a' \wedge \bigwedge_{b\in B,\, 0\neq b\leq a'} \Diamond b \wedge \Box i' \leq x,$$

so $(a', i') \notin f(x)$. By (XI) and Definition A.1.4, we have $(a, i)R(a', i')$, so we conclude that $(a, i) \notin \Box f(x)$. $\Box$

In the proof, we use the completeness of the ortholattice in only two places: first, so that the meets $\bigwedge\{\Diamond b \mid b \in B,\, 0 \neq b \leq a\}$ exist in L for each $a \in B$, and second, in the proof that the map f is surjective. Thus, even if we drop the assumption that L is complete, we can prove the following.

Theorem A.5. Let L be an epistemic ortho-Boolean lattice with distinguished Boolean subalgebra B, satisfying all the conditions of Definition A.1 except possibly completeness. If the meet $\bigwedge\{\Diamond b \mid b \in B,\, 0 \neq b \leq a\}$ exist in L for each $a \in B$, then there is a complete embedding[37] of L into $O(B^{\mathsf{e}})$.

Of course if B is finite, then all the relevant meets exist in L.

Corollary A.6. Let L be an epistemic ortho-Boolean lattice with distinguished *finite* Boolean subalgebra B, satisfying all the conditions of Definition A.1 except possibly completeness. Then there is a complete embedding of L into $O(B^{\mathsf{e}})$.

[37] I.e., an embedding preserving all existing meets and joins.

A.2 Decidability of the logic of epistemic frames of Boolean algebras

In this section, we show that consequence over epistemic frames of Boolean algebras is decidable. The key device is a canonical finite model.

Definition A.7. For any finite $P \subseteq \mathsf{Boole}$, a *$P$-state description* is a conjunction of exactly one of $\mathtt{p}$ or $\neg\mathtt{p}$ for each $\mathtt{p} \in P$. Let W_P be the set of all P-state descriptions, V_P the valuation sending each $\mathtt{p} \in P$ to the set of P-state descriptions in which $\mathtt{p}$ occurs positively, and $\mathcal{M}_P$ the epistemic model of $(\wp(W_P), V_P)$.

For $P \subseteq \mathsf{Boole}$, let $\mathcal{EL}^*(P)$ be the fragment of $\mathcal{EL}^+$ whose only proposition letters belong to P. Let $\mathbf{EB}$ be the class of epistemic frames of Boolean algebras, and define $\vDash_{\mathbf{EB}}$ as in Definition 4.44.

Proposition A.8. For $\varphi, \psi \in \mathcal{EL}^*(P)$, if $\varphi \nvDash_{\mathbf{EB}} \psi$, then there is a possibility in $\mathcal{M}_P$ satisfying φ but not ψ.

Proof. We outline the main steps, leaving some details to the reader. Suppose $\varphi \nvDash_{\mathbf{EB}} \psi$, so there is some Boolean algebra B with valuation θ such that in the epistemic model $\mathcal{M}$ of (B, θ), we have $[\![\varphi]\!]^{\mathcal{M}} \not\subseteq [\![\psi]\!]^{\mathcal{M}}$. Let W_0 be the set of ultrafilters of B, V_0 the valuation for P such that $V(\mathtt{p}) = \{u \in W_0 \mid \theta(\mathtt{p}) \in u\}$, and $\mathcal{M}_0$ the epistemic model of $(\wp(W_0), V_0)$. Then since B embeds into $\wp(W_0)$, $O(B^{\mathsf{e}})$ embeds into $O(\wp(W_0)^{\mathsf{e}})$, from which it is easy to show that $[\![\varphi]\!]^{\mathcal{M}} \not\subseteq [\![\psi]\!]^{\mathcal{M}}$ implies $[\![\varphi]\!]^{\mathcal{M}_0} \not\subseteq [\![\psi]\!]^{\mathcal{M}_0}$. Now let $\sim$ be an equivalence relation on W_0 defined by $u \sim u'$ if $u \cap \{\theta(\mathtt{p}) \mid \mathtt{p} \in P\} = u' \cap \{\theta(\mathtt{p}) \mid \mathtt{p} \in P\}$,[38] and let $[u]$ be the equivalence class of u. Let W_1 be the set of equivalence classes of $\sim$, V_1 the valuation for P such that $V_0(\mathtt{p}) = \{[u] \in W_1 \mid \theta(\mathtt{p}) \in u\}$, and $\mathcal{M}_1$ the epistemic model of $(\wp(W_1), V_1)$. Then one can prove by induction that for any $\chi \in \mathcal{EL}^*(P)$, a possibility (A, I) in $\mathcal{M}_0$ satisfies χ iff the possibility $(\{[u] \mid u \in A\}, \{[u] \mid u \in I\})$ in $\mathcal{M}_1$ satisfies χ. Hence $[\![\varphi]\!]^{\mathcal{M}_0} \not\subseteq [\![\psi]\!]^{\mathcal{M}_0}$ implies $[\![\varphi]\!]^{\mathcal{M}_1} \not\subseteq [\![\psi]\!]^{\mathcal{M}_1}$. Finally, W_1 can be identified with a subset of W_P: send an equivalence class $[u]$ of ultrafilters to the P-state description s such that $\mathtt{p}$ occurs positively in s iff $\theta(\mathtt{p}) \in u$. Then $\wp(W_1)$ embeds into $\wp(W_P)$, so $O(\wp(W_1)^{\mathsf{e}})$ embeds into $O(\wp(W_P)^{\mathsf{e}})$, from which it is easy to show that $[\![\varphi]\!]^{\mathcal{M}_1} \not\subseteq [\![\psi]\!]^{\mathcal{M}_1}$ implies $[\![\varphi]\!]^{\mathcal{M}_P} \not\subseteq [\![\psi]\!]^{\mathcal{M}_P}$. □

For the following corollary, let $\mathcal{EL}^* = \mathcal{EL}^*(\mathsf{Boole})$.

Corollary A.9. For $\varphi, \psi \in \mathcal{EL}^*$, it is decidable whether $\varphi \vDash_{\mathbf{EB}} \psi$.

Proof. Where P is the set of proposition letters appearing in φ or ψ, by Proposition A.8 we can simply check every possibility in the finite model $\mathcal{M}_P$. □

A.3 Normal forms

Closely related to the previous section is a semantic normal form theorem. Given a Boolean formula $\alpha \in \mathcal{EL}^*(P)$, let $st_P(\alpha)$ be the set of all P-state descriptions that entail α in classical logic.

Definition A.10. For any finite $P \subseteq \mathsf{Boole}$ and $\varphi \in \mathcal{EL}(P)$, we say that φ is in *P-canonical normal form* if it is a disjunction of conjunctions of the form

$$\alpha \wedge \bigwedge_{\beta \in st_P(\alpha)} \Diamond\beta \wedge \Box\gamma,$$

where α and γ are disjunctions of P-state descriptions.

[38] Cf. the method of *filtration* in modal logic.

Proposition A.11. For any finite $P \subseteq \mathsf{Boole}$ and $\varphi \in \mathcal{EL}^*(P)$, φ is equivalent over epistemic frames of Boolean algebras to a formula in P-canonical normal form, which can be effectively computed from φ.

Proof. Given a possibility (A, I) in $\mathcal{M}_P$, let α_A be the disjunction of P-state descriptions in A and γ_I the disjunction of P-state descriptions in I, respectively. Then by Lemma 5.4, the formula $\alpha_A \wedge \bigwedge_{\beta \in st_P(\alpha_A)} \Diamond\beta \wedge \Box\gamma_I$ is true only at (A, I) in $\mathcal{M}_P$. It follows that

$$\begin{aligned} [\![\varphi]\!]^{\mathcal{M}_P} &= \bigcup_{(A,I)\in[\![\varphi]\!]^{\mathcal{M}_P}} [\![\alpha_A \wedge \bigwedge_{\beta \in st_P(\alpha_A)} \Diamond\beta \wedge \Box\gamma_I]\!]^{\mathcal{M}_P} \\ &= \bigvee_{(A,I)\in[\![\varphi]\!]^{\mathcal{M}_P}} [\![\alpha_A \wedge \bigwedge_{\beta \in st_P(\alpha_A)} \Diamond\beta \wedge \Box\gamma_I]\!]^{\mathcal{M}_P} \quad \text{since } [\![\varphi]\!]^{\mathcal{M}_P} \text{ is } \between\text{-regular} \\ &= [\![\bigvee_{(A,I)\in[\![\varphi]\!]^{\mathcal{M}_P}} (\alpha_A \wedge \bigwedge_{\beta \in st_P(\alpha_A)} \Diamond\beta \wedge \Box\gamma_I)]\!]^{\mathcal{M}_P}. \end{aligned}$$

Then we claim φ is equivalent to the formula in normal form in the last line. For if not, then by Lemma A.8, there is a possibility in $\mathcal{M}_P$ in which only one of the formulas is true, contradicting the equation above. Clearly the normal form of φ can be effectively computed via the construction of the finite model $\mathcal{M}_P$. $\Box$

From Proposition A.11, we obtain a somewhat unilluminating but nonetheless complete logic with respect to $\vDash_{\mathbf{EB}}$. The key principles beyond orthologic are $\varphi \vdash NF_P(\varphi)$ and $NF_P(\varphi) \vdash \varphi$, where P contains all proposition letters in P and $NF_P(\varphi)$ is the P-canonical normal form of φ computed as in the proof of Proposition A.11 (fixing some ordering of P and W_P to get uniqueness of $NF_P(\varphi)$). Then if $\varphi \vDash_{\mathbf{EB}} \psi$, the proof of Proposition A.11 shows that $NF_P(\psi)$ can be obtained from $NF_P(\varphi)$ by disjunction introduction, so $\varphi \vdash NF_P(\varphi) \vdash NF_P(\psi) \vdash \psi$. We leave it to future work to give a more illuminating set of logical principles that allow us to derive the equivalence between φ and $NF_P(\varphi)$.

References

Maria Aloni. Conceptual covers in dynamic semantics. In Lawrence Cavedon, Patrick Blackburn, Nick Braisby, and Atsushi Shimojima, editors, *Logic, Language and Computation*, volume III. CSLI, 2000.

Maria Aloni. FC disjunction in state-based semantics. Slides for Logical Aspects of Computational Linguistics (*LACL*), Nancy, France, 2018.

Maria Aloni, Luca Incurvati, and Julian J. Schlöder. Epistemic modals in hypothetical reasoning. *Erkenntnis*, 2022. doi: 0.1007/s10670-022-00517-x.

Bob Beddor and Simon Goldstein. Mighty knowledge. *Journal of Philosophy*, 118(5):229–269, 2021. doi: 10.5840/jphil2021118518.

J. van Benthem. *Exploring Logical Dynamics*. Center for the Study of Language and Information, Stanford, CA, 1996.

Johan van Benthem, Nick Bezhanishvili, and Wesley H. Holliday. A bimodal perspective on possibility semantics. *Journal of Logic and Computation*, 27(5):1353–1389, 2017. doi: 10.1093/logcom/exw024.

Katalin Bimbó. Functorial duality for ortholattices and De Morgan lattices. *Logica Universalis*, 1:311–333, 2007. doi: 10.1007/s11787-007-0016-9.

G. Birkhoff. *Lattice Theory.* American Mathematical Society, New York, 1940.

G. Birkhoff. *Lattice Theory.* American Mathematical Society, Providence, RI, 3rd edition, 1967.

Garrett Birkhoff and John von Neumann. The logic of quantum mechanics. *Annals of Mathematics*, 37(4): 823–843, 1936.

P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic.* Cambridge University Press, New York, 2001.

Justin Bledin. Logic informed. *Mind*, 123(490):277–316, 2014. doi: 10.1093/mind/fzu073.

Justin Bledin. Modus Ponens defended. *The Journal of Philosophy*, CXII(2):57–83, 2015. doi: 10.5840/jphil201511225.

John P. Burgess. Which modal models are the right ones (for logical necessity)? *Theoria*, 18(47):145–158, 2003. doi: 10.1387/theoria.418.

A. Chagrov and M. Zakharyaschev. *Modal logic*, volume 35 of *Oxford Logic Guides*. The Clarendon Press, New York, 1997.

Nate Charlow. Grading modal judgement. *Mind*, 129(515):769–807, 2019. doi: 10.1093/mind/fzz028.

Maria Luisa Dalla Chiara and Roberto Giuntini. Quantum logics. In Dov Gabbay and Franz Guenthner, editors, *Handbook of Philosophical Logic*, pages 129–228. Springer, 2002.

H. Dishkant. Semantics of the minimal logic of quantum mechanics. *Studia Logica*, 30(4):23–30, 1972. doi: 10.1007/BF02120818.

H. Dishkant. Imbedding of the quantum logic in the modal system of Brouwer. *The Journal of Symbolic Logic*, 42(3):321–328, 1977. doi: 10.2307/2272861.

Anna Dmitrieva. Positive modal logic beyond distributivity: duality, preservation and completeness. Master's thesis, University of Amsterdam, 2021.

Jaap van der Does, Willem Groeneveld, and Frank Veltman. An update on "might". *Journal of Logic, Language, and Information*, 6:361–381, 1997.

Cian Dorr and John Hawthorne. Embedding epistemic modals. *Mind*, 122(488):867–913, 2013. doi: 10.1093/mind/fzt091.

M. A. E. Dummett and E. J. Lemmon. Modal logics between S4 and S5. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 5:250–264, 1959. doi: 10.1002/malq.19590051405.

Kai von Fintel and Anthony Gillies. Must...stay...strong! *Natural Language Semantics*, 18(4):351–383, 2010. doi: 10.1007/s11050-010-9058-2.

Kai von Fintel and Anthony Gillies. Still going strong. *Natural Language Semantics*, 29:91–113, 2021. doi: 10.1007/s11050-020-09171-x.

Frederic B. Fitch. *Symbolic Logic: An Introduction.* The Ronald Press Company, New York, 1952.

Frederic B. Fitch. Natural deduction rules for obligation. *American Philosophical Quarterly*, 3(1):27–38, 1966.

Vera Flocke. Nested epistemic modals. Manuscript, Indiana University, Bloomington, July 2020.

Bas van Fraassen. Probabilities of conditionals. In Harper and Hooker, editors, *Foundations of Probability Theory, Statistical Inference, and Statistical Theories of Science*, volume I, pages 261–308. D. Reidel Publishing Company, Dordrecht-Holland, 1976.

Christopher Gauker. A strictly stronger relative *must*. *Thought: A Journal of Philosophy*, 10:82–89, 2020. doi: 10.1002/tht3.482.

Anthony S. Gillies. Updating data semantics. *Mind*, 129(513):1–41, 2020. doi: 10.1093/mind/fzy008.

Kurt Gödel. Eine Interpretation des intuitionistischen aussagenkalküls. *Ergebnisse eines Mathematischen Kolloquiums*, 4:39–40, 1933.

Robert I. Goldblatt. Semantic analysis of orthologic. *Journal of Philosophical Logic*, 3(1):19–35, 1974. doi: 10.1007/BF00652069.

Robert I. Goldblatt. The Stone space of an ortholattice. *Bulletin of the London Mathematical Society*, 7(1): 45–48, 1975. doi: 10.1112/blms/7.1.45.

Simon Goldstein. Generalized update semantics. *Mind*, 128(511):795–835, 2019. doi: 10.1093/mind/fzy076.

Simon Goldstein. Epistemic modal credence. *Philosophers' Imprint*, 26, 2021.

Jeroen Groenendijk, Martin Stokhof, and Frank Veltman. Coreference and modality. In Shalom Lappin, editor, *Handbook of Contemporary Semantic Theory*, pages 179–216. Oxford, Blackwell, 1996.

A. Grzegorczyk. A philosophically plausible formal interpretation of intuitionistic logic. *Indagationes Mathematicae*, 26:596–601, 1964.

Ian Hacking. Possibility. *Philosophical Review*, 76(2):143–168, 1967.

John Harding. Varieties of ortholattices containing the orthomodular lattices. Master's thesis, McMaster University, 1988.

Peter Hawke and Shane Steinert-Threlkeld. Semantic expressivism for epistemic modals. *Linguistics and Philosophy*, 44:475–511, 2021. doi: 10.1007/s10988-020-09295-7.

Jaakko Hintikka. *Knowledge and Belief: An Introduction to the Logic of the Two Notions*. Cornell University Press, Ithaca, NY, 1962.

Wesley H. Holliday. Partiality and adjointness in modal logic. In R. Goré, B. Kooi, and A. Kurucz, editors, *Advances in Modal Logic*, volume 10, pages 313–332. College Publications, London, 2014.

Wesley H. Holliday. Three roads to complete lattices: Orders, compatibility, polarity. *Algebra Universalis*, 82(26), 2021a. doi: 10.1007/s00012-021-00711-y.

Wesley H. Holliday. Possibility semantics. In Melvin Fitting, editor, *Selected Topics from Contemporary Logics*, Landscapes in Logic, pages 363–476. College Publications, 2021b. URL https://escholarship.org/uc/item/9ts1b228.

Wesley H. Holliday. Compatibility and accessibility: lattice representations for semantics of non-classical and modal logics. In David Fernández Duque and Alessandra Palmigiano, editors, *Advances in Modal Logic, Vol. 14*, pages 507–529. London, 2022. arXiv:2201.07098 [math.LO].

Wesley H. Holliday. Possibility frames and forcing for modal logic. *The Australasian Journal of Logic*, Forthcoming. URL `https://escholarship.org/uc/item/0tm6b30q`.

Wesley H. Holliday and Thomas F. Icard. Moorean Phenomena in Epistemic Logic. In Lev Beklemishev, Valentin Goranko, and Valentin Shehtman, editors, *Advances in Modal Logic*, volume 8, pages 178–199. College Publications, London, 2010.

Wesley H. Holliday and Thomas F. Icard. Axiomatization in the meaning sciences. In Derek Ball and Brian Rabern, editors, *The Science of Meaning: Essays on the Metatheory of Natural Language Semantics*, pages 73–97, Oxford, 2018. Oxford University Press. doi: 10.1093/oso/9780198739548.003.0002.

Wesley H. Holliday and Tadeusz Litak. Complete additivity and modal incompleteness. *The Review of Symbolic Logic*, 12(3):487–535, 2019. doi: 10.1017/S1755020317000259.

Wesley H. Holliday and Matthew Mandelkern. Compatibility, compossibility, and epistemic modality. In Marco Degano, Tom Roberts, Giorgio Sbardolini, and Marieke Schouwstra, editors, *Proceedings of the 23rd Amsterdam Colloquium*, pages 120–126, 2022.

Wesley H. Holliday, Tomohiro Hoshi, and Thomas F. Icard. A uniform logic of information dynamics. In Thomas Bolander, Torben Braüner, Silvio Ghilardi, and Lawrence Moss, editors, *Advances in Modal Logic*, volume 9, pages 348–367. College Publications, London, 2012.

Wesley H. Holliday, Tomohiro Hoshi, and Thomas F. Icard. Information dynamics and uniform substitution. *Synthese*, 190(1):31–55, 2013. doi: 10.1007/s11229-013-0278-0.

L. Humberstone. From worlds to possibilities. *Journal of Philosophical Logic*, 10(3):313–339, 1981. doi: 10.1007/BF00293423.

Luca Incurvati and Julian J. Schlöder. Epistemic multilateral logic. *Review of Symbolic Logic*, 2020. doi: 10.1017/S1755020320000313.

Justin Khoo. *The Meaning of* If. Oxford University Press, 2022.

Justin Khoo and Paolo Santorio. Lecture notes: Probability of conditionals in modal semantics. Unpublished manuscript, July 2018.

Nathan Klinedinst and Daniel Rothschild. Connectives without truth-tables. *Natural Language Semantics*, 20:137–175, 2012. doi: 10.1007/s11050-011-9079-5.

M. Kolibiar. Distributive sublattices of a lattice. *Proceedings of the American Mathematical Society*, 34(2): 359–364, 1972.

Niko Kolodny and John MacFarlane. Ifs and oughts. *Journal of Philosophy*, 107(3):115–143, 2010. doi: 10.5840/jphil2010107310.

Angelika Kratzer. What 'must' and 'can' must and can mean. *Linguistics and Philosophy*, 1(3):337–355, 1977. doi: 10.1007/BF00353453.

Angelika Kratzer. The notional category of modality. In H. Eikmeyer and H. Rieser, editors, *Words, Worlds, and Contexts: New Approaches in Word Semantics*, pages 38–74. de Gruyter, 1981.

Angelika Kratzer. An investigation of the lumps of thought. *Linguistics and Philosophy*, 12(5):607–653, 1989. doi: 10.1007/BF00627775.

Angelika Kratzer. What's an epistemic modal anyway? Manuscript, University of Massachusetts at Amherst, July 2020.

Saul A. Kripke. Semantical analysis of intuitionistic logic I. In J. N. Crossley and M. A. E. Dummett, editors, *Formal Systems and Recursive Functions*, pages 92–130. North-Holland, Amsterdam, 1965.

Daniel Lassiter. *Must*, knowledge, and (in)directness. *Natural Language Semantics*, 24(2):117–163, 2016.

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

John MacFarlane. Epistemic modals are assessment sensitive. In Andy Egan and Brian Weatherson, editors, *Epistemic Modality*, pages 144–177. Oxford University Press, 2011.

M. Donald MacLaren. Atomic orthocomplemented lattices. *Pacific Journal of Mathematics*, 14(2):597–612, 1964.

Matthew Mandelkern. Bounded modality. *The Philosophical Review*, 128(1):1–61, 2019. doi: 10.1215/00318108-7213001.

Matthew Mandelkern. Dynamic non-classicality. *Australasian Journal of Philosophy*, 98(2):382–392, 2020. doi: 10.1080/00048402.2019.1623826.

Matthew Mandelkern. Bounds: Meaning and the limits of interpretation. Manuscript, NYU, January 2023.

Joseph McDonald and Kentarô Yamamoto. Choice-free duality for orthocomplemented lattices. arXiv:2010.06763 [math.LO], 2021.

Vann McGee. A counterexample to modus ponens. *The Journal of Philosophy*, 82(9):462–471, 1985. doi: 10.2307/2026276.

J. C. C. McKinsey and Alfred Tarski. Some theorems about the sentential calculi of Lewis and Heyting. *The Journal of Symbolic Logic*, 13:1–15, 1948. doi: 10.2307/2268135.

Sarah Moss. On the semantics and pragmatics of epistemic vocabulary. *Semantics and Pragmatics*, 8(5): 1–81, 2015. doi: 10.3765/sp.8.5.

Dilip Ninan. Quantification and epistemic modality. *The Philosophical Review*, 127(2):433–485, 2018. doi: 10.1215/00318108-6973010.

Kurt Norlin. In the logic of certainty, $\supset$ is stronger than $\to$. *Thought*, 9(1):58–63, 2020.

Kurt Norlin. Acceptance and certainty, doxastic modals, and indicative conditionals. *Journal of Philosophical Logic*, 2022. doi: 10.1007/s10992-022-09656-6.

Daniel Rothschild. Expressing credences. *Proceedings of the Aristotelian Society*, 112(1):99–114, 2011. doi: 10.1111/j.1467-9264.2012.00327.x.

Paolo Santorio. Trivializing informational consequence. *Philosophy and Phenomenological Research*, 2021. doi: 10.1111/phpr.12745.

Robert Stalnaker. Pragmatic presuppositions. In Milton K. Munitz and Peter Unger, editors, *Semantics and Philosophy*, pages 197–213. New York University Press, New York, 1974. doi: 10.1093/0198237073.003.0003.

Robert C. Stalnaker. A theory of conditionals. In Nicholas Rescher, editor, *Studies in Logical Theory*, pages 98–112. Blackwell, 1968.

Robert C. Stalnaker. Probability and conditionals. *Philosophy of Science*, 37(1):64–80, 1970.

Una Stojnić. One's *Modus Ponens*: Modality, coherence, and logic. *Philosophy and Phenomenological Research*, 93(3):1–48, 2016. doi: 10.1111/phpr.12307.

Eric Swanson. The application of constraint semantics to the language of subjective uncertainty. *Journal of Philosophical Logic*, 45(2):121–146, 2016. doi: 10.1007/s10992-015-9367-5.

Alfred Tarski. Zur grundlegung der bool'schen algebra. I. *Fundamenta Mathematicae*, 24:177–198, 1935.

Josef Tkadlec. Partially additive states on orthomodular posets. *Colloquium Mathematicum*, 62:7–14, 1991.

Josef Tkadlec. Partially additive measures and set representations of orthoposets. *Journal of Pure and Applied Algebra*, 86:79–94, 1993.

Janneke van Wijnbergen-Huitink. Modal concord. In Daniel Gutzmann, Lisa Matthewson, Cécile Meier, Hotze Rullman, and Thomas Ede Zimmermann, editors, *The Wiley Blackwell Companion to Semantics*. Wiley Online Library, 2020. doi: 10.1002/9781118788516.sem129.

Frank Veltman. *Logics for Conditionals*. PhD thesis, University of Amsterdam, 1985.

Frank Veltman. Defaults in update semantics. *Journal of Philosophical Logic*, 25(3):221–261, 1996. doi: 10.1007/BF00248150.

Malte Willer. Dynamics of epistemic modality. *Philosophical Review*, 122(1):45–92, 2013. doi: 10.1215/00318108-1728714.

Ludwig Wittgenstein. *Philosophical Investigations*. Wiley-Blackwell, 2001, 3rd edition, 1953.

Seth Yalcin. Epistemic modals. *Mind*, 116(464):983–1026, 2007. doi: 10.1093/mind/fzm983.

Seth Yalcin. Dynamic semantics. In Gillian Russell and Delia Graff Fara, editors, *Routledge Companion to Philosophy of Language*, pages 253–279. Routledge, New York, 2012a.

Seth Yalcin. Context probabilism. In Maria Aloni, Vadim Kimmelman, Floris Roelofsen, Galit W. Sasson, Katrin Schulz, and Matthijs Westera, editors, *Logic, Language, and Meaning. 18th Amsterdam Colloquium*, volume 7218 of *Lecture Notes in Computer Science*, pages 12–21, 2012b.

Seth Yalcin. A counterexample to modus tollens. *Journal of Philosophical Logic*, 41:1001–1024, 2012c. doi: 10.1007/s10992-012-9228-4.

Seth Yalcin. Epistemic modality *de re*. *Ergo*, 2(19):475–527, 2015. doi: 10.3998/ergo.12405314.0002.019.