

BELIEVING EPISTEMIC CONTRADICTIONS

Abstract

What is it to believe something might be the case? We develop a puzzle that creates difficulties for standard answers to this question. We go on to propose our own solution, which integrates a Bayesian approach to belief with a dynamic semantics for epistemic modals. After showing how our account solves the puzzle, we explore a surprising consequence: virtually all of our beliefs about what might be the case provide counterexamples to the view that rational belief is closed under logical implication.

1 The Puzzle

Ari the burglar has been casing the house for hours. As far as she can tell, not a mouse is stirring. Consequently, Ari believes the house is empty.

Still, Ari is an experienced burglar; she knows that even the most thorough reconnaissance is fallible. Thus she allows that there's some possibility—albeit very remote—that an inconspicuous resident is still inside.

Given this, it seems we should be able to report Ari's belief state as follows:

- (1) # Ari believes the house is empty and might not be empty.

But (1) sounds odd. In uttering (1), it sounds like you're attributing to Ari an incoherent belief.¹ This is surprising, since nothing in our set-up seemed to saddle Ari with incoherence. On the face of it, Ari is being perfectly rational: forming beliefs in response to her evidence, without being overly confident in her conclusions.

We have here the makings of a puzzle, which can be formulated in more general terms. It's a sad but indisputable fact that humans are fallible: we've all held beliefs that turned out to be mistaken. Recognizing this fact, it seems rational for an agent to hold some belief, while also acknowledging that it's possible that this belief is mistaken. At the very least, it seems that this stance is *coherent*. Letting ' \diamond ' represent epistemic possibility, we can formulate this principle as follows:

¹It's important that both conjuncts of the complement clause (*the house is empty and (it) might not be empty*) scope under *believes*. The sentence:

- (i) Ari believes [the house is empty]; and it might not be empty.

does not ascribe to Ari an incoherent belief. We discuss versions of (i) in §3.

Fallibility It's sometimes coherent for an agent to believe ϕ and also believe $\diamond\neg\phi$.

For example, Ari the burglar believes that the house is empty, while also believing that there's some possibility that someone is still inside.

While Fallibility seems plausible, our reaction to (1) motivates a principle that stands in tension with it. Following Yalcin (2007), let us refer to any sentence of the form, $\phi \wedge \diamond\neg\phi$ as an *epistemic contradiction*. Given the oddity of (1), it's tempting to maintain that it is always incoherent to believe an epistemic contradiction:

No Contradictions It's never coherent to believe $(\phi \wedge \diamond\neg\phi)$.

But it is *prima facie* difficult to reconcile these two principles. Is any story about beliefs involving epistemic modals up to the task?

2 Roadmap

Whenever we encounter two principles that stand in tension with each other, it is prudent to examine whether the principles are really all that plausible. We thus start by defending our principles: in §3, we argue that both are well-supported by linguistic data. Consequently we should solve our puzzle not by denying one of our principles, but rather by providing a semantics for epistemic modals and belief reports that validates both of them, thereby dissolving the appearance of inconsistency.

This is no trivial task. The classical semantics for epistemic modals doesn't deliver the desired results: as we show in §4, standard contextualist (Kratzer 1981, 1991, 2012) and relativist (Egan 2007; Stephenson 2007a,b; MacFarlane 2011) approaches to modals are forced to give up either Fallibility or No Contradictions.

By itself, this may strike some as unsurprising. After all, Yalcin (2007, 2011) has argued that the classical semantics has trouble explaining the infelicity of epistemic contradictions in embedded contexts. However, we show that while leading non-classical accounts of beliefs involving epistemic modals (Veltman 1996; Gillies 2001; Yalcin 2007, 2011, 2012a; Willer 2013) validate No Contradictions, they invalidate Fallibility (§5). As a result, they also fail to provide a satisfactory solution to our puzzle.

In §6 we present our own solution, which integrates a Bayesian approach to belief with a dynamic semantics for epistemic modals. According to our theory, A believes ϕ iff A assigns a sufficiently high credence to the result of updating A's information with ϕ . Assuming a dynamic account of updating, this theory validates both of our principles, thereby solving the puzzle. In broad outline, the solution is this. Our theory borrows from Bayesianism the insight that belief requires sufficiently high credence, not certainty. At the same time, the dynamic account of updating predicts that modal beliefs are 'transparent': to believe that the house might not be empty is to be less than certain that it's

empty. And so our theory predicts that whenever an agent believes ϕ without being certain of it, she counts as believing $\Diamond\neg\phi$, thereby validating Fallibility. However, the dynamic account of updating also predicts that no body of information can coherently be updated with $\phi \wedge \Diamond\neg\phi$. Hence no coherent agent will believe an epistemic contradiction, thereby validating No Contradictions.

After developing our solution to the puzzle, we defend and extend it. §7 explores a surprising consequence of our account of belief: coherent agents can believe ϕ and believe $\Diamond\neg\phi$, but they cannot coherently believe an obvious entailment of these beliefs: $\phi \wedge \Diamond\neg\phi$. Thus our solution entails that rational belief is not closed under logical implication. Indeed, our view suggests that counterexamples to closure occur all the time—no need to look to lotteries and prefaces. Finally, §8 defends our solution from objections.

3 Defending our Principles

3.1 Fallibility

Let us start with Fallibility—the thesis that it’s sometimes coherent to believe ϕ and to also believe $\Diamond\neg\phi$. In addition to its intuitive plausibility, we offer three arguments in support of this thesis.

The first is the argument from *concessive belief attributions*. Let a concessive belief attribution (CBA) be a discourse of the form:

- (2) I believe ϕ . But $\Diamond\neg\phi$.

Such discourses seem perfectly felicitous. Consider, for instance:

- (3) I believe the movie starts at 7, but $\left\{ \begin{array}{l} \text{it might start later} \\ \text{I might be mistaken} \end{array} \right\}$.²

Presumably it is infelicitous to make an assertion if it is impossible for a coherent agent to have a true belief in its content. But if Fallibility were false, no coherent agent could truly believe a CBA. After all, suppose the first conjunct of a CBA (*I believe ϕ*) is true. If Fallibility were false, the second conjunct (*But $\Diamond\neg\phi$*) could not be coherently believed. More generally, if Fallibility were false, then anyone who uttered a CBA would be committed to having any incoherent doxastic state.

The second argument for Fallibility comes from considering discourses of the form:

- (4) Ari believes the house is empty. But she realizes/recognizes that it might not be.

²Concessive knowledge attributions (that is, sentences of the form, *I know ϕ , but $\Diamond\neg\phi$*) have attracted a great deal of attention in recent years. Interestingly, while it’s frequently observed that concessive knowledge attributions are infelicitous (see, e.g., Lewis 1996; Rysiew 2001; Stanley 2005; Dodd 2011; Worsnip 2015), the comparative felicity of CBAs seems to have gone unnoticed.

(4) sounds perfectly fine.³ But presumably if an agent realizes or recognizes ϕ , it follows that she believes ϕ . So if Fallibility were false, (4) would be incoherent.

Our third argument for Fallibility is more theoretical. According to this argument, Fallibility follows from two further principles, both of which are supported by independent data. The first principle is that it's coherent to hold uncertain beliefs:

Uncertain Belief It's possible to coherently believe ϕ without being certain that ϕ .

The idea that belief doesn't require certainty is widely taken for granted.⁴ It's also well-supported by our everyday *belief* and *certainty* talk. Some examples:

- (5) I believe, but am not certain, that the land was granted to Mr. Baca for pasturing purposes...⁵
- (6) I believe but am not certain that Palestine is not a party to the statute of the International Court of Justice, and I believe but am not certain that the General Assembly has not addressed an invitation to Palestine to join the Convention.⁶

Such utterances seem to express coherent doxastic states.⁷

While (5) and (6) are first-person attitude ascriptions, it is also natural to ascribe uncertain beliefs to others. For example:

- (7) Ari believes that the house is empty, but she's not certain of it.

If Uncertain Belief were false, (7) would ascribe to Ari an incoherent doxastic state. But to our ears, (7) seems perfectly coherent.

The second principle posits a connection between uncertainty and believing possible:

Uncertainty-Possibility Link If an agent A is coherent, then if A isn't certain that ϕ , A is in a position to believe $\diamond\neg\phi$.

We suspect that Uncertainty-Possibility Link will prove more controversial than Uncertain Belief. A couple of clarificatory remarks may help forestall some immediate objections. First, in order for the principle to be remotely plausible, it's important that it is interpreted as making a claim about *epistemic* possibility, rather than physical or metaphysical possibility. Second, the 'in a position' qualification here is important. To see this, consider an agent who is not certain

³In a similar vein, Hawthorne et al. observe that it sounds coherent to say, *I believe it's raining, but I know it might not be* (2016: 1396).

⁴Though see Clarke (2013) and Dodd (forthcoming) for dissent.

⁵*Executive Documents of the House of Representatives*, vol.125-126: 40.

⁶<https://lettersblogatory.com/2015/01/02/palestine-signs-new-york-convention/>

⁷See Christensen (2005: 21) and Hawthorne et al. (2016: 1395) for similar observations.

that ϕ , but has never even considered the question of whether ϕ . Arguably, it is at least somewhat counterintuitive to say that they believe that ϕ might be false. Thus, Uncertainty-Possibility Link should be interpreted as making the following claim: if a coherent agent isn't certain that ϕ , they are committed to believing (perhaps on further reflection) that it's epistemically possible that ϕ might be false.

Even with these clarifications, some may deem Uncertainty-Possibility Link implausible. After all, it's widely held that certainty is a very demanding state. For example, many philosophers maintain that if one is certain of ϕ , one should be willing to accept a bet where one wins a penny if ϕ is true, and one loses one's life otherwise. Presumably, very few of ordinary beliefs are held with this degree of conviction.⁸ But if for virtually any ordinary proposition ϕ , we do not believe ϕ with certainty, Uncertainty-Possibility Link entails that we are almost always in a position to believe $\diamond\neg\phi$, and hence almost never in a position to believe $\neg\diamond\neg\phi$.⁹ But this seems wrong. In ordinary contexts, it would be natural for a speaker to say:

(8) Jim believes there's no possibility the Lions will win.

even if she doesn't think that Jim would bet his life on the claim that the Lions will lose.

However, proponents of Uncertainty-Possibility Link have a natural response. Observe that in ordinary contexts, it would be equally natural for a speaker to say:

(9) Jim is certain that the Lions will lose.

This suggests that truth-values of certainty ascriptions differ from context to context, where here the context could be either that of the speaker or that of the subject of ascription.¹⁰ According to this view, only in the most demanding contexts will a claim of the form, *A is certain that ϕ* entail that A is willing to bet their life on ϕ . In most ordinary contexts, less is required.¹¹

Indeed, reflecting on the assertability conditions of certainty ascriptions provides a compelling argument for Uncertainty-Possibility Link. Consider:

⁸For the view that certainty is rarely (if ever) attained, see Russell (1912); Unger (1975), among others.

⁹Given the standard assumption that *must* (\Box) and *might* are duals (i.e., $\Box\phi$ iff $\neg\diamond\neg\phi$), this is equivalent to saying that we are almost never in a position to believe $\Box\phi$.

¹⁰See Lewis (1979: 353-354) and Stanley (2008). Note that if certainty ascriptions are context-sensitive, then it would be more accurate to formulate Uncertainty-Possibility Link in the formal mode: if *A isn't certain that ϕ* is true in a context, then *A is in a position to believe $\diamond\neg\phi$* will also be true in that context.

¹¹Some authors will resist the conclusion that certainty attributions are context-sensitive. These authors will insist that the speaker's utterance of (9) is "loose talk": false but pragmatically acceptable (Unger 1975; Lasersohn 1999; Kennedy 2007.) Positing such rampant falsity in ordinary conversation strikes us as a cost to such a view. That said, the "loose talk" approach is still compatible with Uncertainty-Possibility Link, provided it is reformulated in terms of assertability conditions: if *A isn't certain that ϕ* is assertable in a context, then *A is in a position to believe $\diamond\neg\phi$* will also be assertable in that context.

- (10) a. Ari isn't certain that the house is empty.
 b. ?? But she doesn't believe/think there's any possibility there's someone inside.

Assuming we're in a context in which Ari has considered the question of whether the house is empty, following (10a) with (10b) sounds very odd. Uncertainty-Possibility Link explains this: (10) ascribes an incoherent doxastic state to Ari.¹²

Thus, both Uncertain Belief and Uncertainty-Possibility Link are well-supported by the data. And their conjunction entails Fallibility. After all, by Uncertain Belief it's coherent for an agent to believe ϕ without being certain of ϕ . By Uncertainty-Possibility Link, such an agent will always be in a position to believe $\diamond\neg\phi$. *A fortiori*, it will sometimes be coherent for an agent to believe both ϕ and $\diamond\neg\phi$, as Fallibility maintains.

We have presented three arguments for Fallibility: the argument from CBAs, the argument from the *realize-belief* entailment, and the argument from the conjunction of Uncertain Belief and Uncertainty-Possibility Link. Of course, not all readers will be fully persuaded by these arguments; we will discuss some objections to Fallibility in §8. But for now, we propose to accept Fallibility as a working hypothesis.

3.2 No Contradictions

Turn now to No Contradictions. We've already provided the main argument for this principle: (1) (*Ari believes the house is empty and might not be empty*) seems to ascribe an incoherent doxastic state to Ari. What's more, nothing hinges on the details of the example. As far as we can tell, every instance of the schema:

- (11) # A believes $(\phi \wedge \diamond\neg\phi)$.

seems to ascribe incoherent beliefs to A.

We can strengthen the case for No Contradictions by observing that the oddity of (11) is not an isolated phenomenon. Epistemic contradictions sound incoherent in a variety of contexts. It's often been noted that assertions of unembedded epistemic contradictions sound bizarre:

- (12) # The house is empty and might not be empty.¹³

Various authors have also noted that epistemic contradictions sound odd in the antecedents of indicative conditionals, as well as under other attitude

¹²One potential complication for this argument is that *believes* is neg-raising, and so perhaps the logical form of *A doesn't believe ϕ* is really, *A believes $\neg\phi$* . To control for this, one can rewrite the sentence with the quantifier, *nobody*. To our ears, *Nobody is certain that the house is empty, but nobody thinks there's any possibility there's someone inside* is equally marked.

¹³See Veltman (1996); Gillies (2001); Yalcin (2007, 2011); Willer (2013).

verbs—for instance, *suppose* and *imagine* (Yalcin 2007, 2011; Anand and Hacquard 2013; Dorr and Hawthorne 2013). This suggests that there is a genuine and general phenomenon here.¹⁴

Some readers may be inclined to concede instances of (11) are infelicitous, but doubt that this is best explained by the hypothesis that (11) ascribes incoherent beliefs to agents. But then what explains the infelicity of (11), if not the incoherence of the underlying doxastic state? Providing an explanation is no easy matter. Indeed, the puzzle can be reframed in a way that relies directly on our linguistic judgments, rather than No Contradictions. What account of beliefs involving epistemic modals will validate Fallibility, while also explaining the infelicity of (11)?

Having motivated our two principles, we now turn to consider whether any account of beliefs involving epistemic modals can validate both of them.

4 Troubles for the Classical Semantics

Orthodoxy has it that modals quantify over possibilities (Kratzer 1981, 1991, 2012). But they don't quantify over just any possibilities. Their domain is restricted by a contextually determined set of worlds: *the modal base*. Possibility modals (e.g., *might*) existentially quantify over the modal base; necessity modals (e.g., *must*) universally quantify over it. Let c be a context, w an index, $\llbracket \phi \rrbracket^c$ the set of indices w such that ϕ is true at $\langle c, w \rangle$, and $B_{c,w}$ the modal base determined by c and w . Contextualists propose:

Contextualist Might $\llbracket \diamond \phi \rrbracket^{c,w} = 1$ iff $B_{c,w} \cap \llbracket \phi \rrbracket^c \neq \emptyset$.

Epistemic modals are evaluated using an *epistemic modal base*: a set of worlds reflecting the epistemic state of some contextually determined agents. Thus:

(13) The house might not be empty.

is true iff the proposition \langle The house is not empty \rangle is compatible with the epistemic state of the contextually determined group.¹⁵

Standard relativist accounts of epistemic modals (Egan 2007, Stephenson 2007a,b; MacFarlane 2011) are similar. Their main point of departure is that they take the epistemic modal base to reflect the epistemic state of an assessor—an individual who is interpreting the modal. Let a be a context of assessment and $B_{c,a}$ the c, a -determined modal base. Relativists propose:

Relativist Might $\llbracket \diamond \phi \rrbracket^{c,w,a} = 1$ iff $B_{c,a} \cap \llbracket \phi \rrbracket^c \neq \emptyset$.

¹⁴Note that both embedded and unembedded epistemic contradictions remain infelicitous when a contrast marker such as *but* is used. For instance, *Ari believes the house is empty but might not be* is quite bizarre (even if it is slightly less bizarre than (1)).

¹⁵Our exposition makes a standard simplification by omitting the role that the ordering source plays in most versions of a Kratzerian semantics. We also omit the fact that some take epistemic modals to convey indirectness (see e.g., von Stechow and Gillies 2010).

This allows that an utterance of (13) could be true relative to one context of assessment (where the relevant folks' epistemic state leaves open the possibility that the house is occupied), and false relative to a different context of assessment (where the relevant folks' epistemic state entails the house is empty).

In order for the classical semantics to generate predictions about our puzzle, we need to say more about the nature of the epistemic modal base. In what follows, we look at the two most natural options: the *knowledge-based approach* and the *belief-based approach*. We argue that neither validates both of our principles; hence neither provides a satisfactory resolution to our puzzle.

According to the knowledge-based approach, an epistemic modal base is the set of possibilities compatible with what the relevant agents know, or can come to know (Hacking 1967; Kratzer 1981, 2012; DeRose 1991; Egan et al. 2005; Stanley 2005; Stephenson 2007a; Hawthorne 2007, 2012; Egan and Weatherson 2011). On this approach, if A believes $(\phi \wedge \Diamond \neg\phi)$, then what A believes is equivalent to:

- (14) ϕ and $(\neg\phi$ is compatible with what the relevant agents know).

But why would this be incoherent? Suppose A is the only relevant agent. Since belief doesn't entail knowledge, A could believe ϕ without knowing ϕ . What's more, A could truly believe she's in such a position—that is, she could truly believe that she believes ϕ without knowing ϕ .¹⁶ Thus the knowledge-based approach fails to predict No Contradictions.

Some might think that this is too quick. Perhaps even though a subject can believe ϕ while failing to know ϕ , no subject can coherently take herself to believe ϕ while failing to know ϕ . One way of motivating this would be to appeal to the idea that knowledge is the norm of belief (Williamson 2000: 47; Sutton 2005, 2007; Bird 2007; Huemer 2007).¹⁷

But we find this strategy unconvincing. There are certainly agents who take themselves to hold beliefs that don't amount to knowledge. Consider Thelma the theist, who professes to believe that God exists, while also claiming that she doesn't know that God exists: it's a matter of faith. Or consider Louise the lottery ticket holder, who believes her lottery ticket will lose (on statistical grounds), but also claims not to know it will lose, on the grounds that knowledge requires safety (McGlynn 2013). It's natural to describe their doxastic states thus:

- (15) ✓ Thelma believes that God exists and that she doesn't know God exists.
 (16) ✓ Louise believes that her ticket will lose and that she doesn't know her ticket will lose.

¹⁶If there are relevant agents besides A, we expect cases in which one can truly believe an epistemic contradiction to be even more common: they'll include any case in which A believes ϕ and also believes that the other relevant folks don't know ϕ .

¹⁷In a similar vein, Dorr and Hawthorne (2013: 910, n.60) suggest that some uses of *believes* implicate that the believer believes that she knows the complement clause.

But the knowledge-based approach predicts that (15) and (16) are equivalent to:

- (17) # Thelma believes God exists and might not exist.
- (18) # Louise believes her ticket will lose and might win.

So the knowledge-based approach does not explain the difference in felicity between these pairs of sentences.¹⁸

Given these difficulties for the knowledge-based approach, one might adopt a belief-based approach, according to which an epistemic modal base is the set of possibilities compatible with what the relevant agents believe. This approach is in a better position to capture No Contradictions. According to the belief-based approach, if A believes $(\phi \wedge \Diamond\neg\phi)$, then what A believes is equivalent to a conjunction of the form:

- (19) ϕ and $(\neg\phi$ is compatible with what the relevant agents believe).

If we assume that the relevant folks typically include A, this entails that A is committed to believing a Moore-paradoxical proposition:

- (20) ϕ and I don't believe ϕ .

It's a familiar observation that such Moorean beliefs seem incoherent. And so the incoherence of believing an epistemic contradiction is explained in terms of the incoherence of Moorean belief.¹⁹

However, giving this explanation of No Contradictions requires giving up Fallibility. Here's why. Suppose Fallibility holds, and hence that there's a coherent agent who believes ϕ and also believes $\Diamond\neg\phi$. For concreteness, let's focus on Ari, who believes the house is empty, but also believes the house might not be empty. Given the belief-based approach, it follows that Ari believes that it's compatible with what she believes that the house isn't empty. And so Ari is committed to having a Moorean belief: \langle The house is empty and I don't believe the house is empty \rangle . But then Ari is just as incoherent as someone who believes an epistemic contradiction. Thus the belief-based approach only vindicates No Contradictions at the expense of giving up Fallibility.

At this point, some may question the way we have argued against the classical semantics. We've been assuming that the epistemic modal base is *either* the set of possibilities compatible with what the relevant agents know or the

¹⁸While our main criticism of the knowledge-based approach is that it doesn't validate No Contradictions, it's also unclear whether it validates Uncertainty-Possibility Link. The only way that the knowledge-based approach could validate this principle is if knowing ϕ entails being in a position to be certain that ϕ . While some authors endorse this entailment (Ayer 1936; Moore 1959; Unger 1975), cases such as Radford's (1966) unconfident examinee provide grounds for doubt. Intuitively, the examinee knows the answer to the examiner's question, even though he isn't in a position to be certain of it (Armstrong 1973; Stanley 2008; McGlynn 2014).

¹⁹Of course, there remains a further question as to why Moorean beliefs are incoherent. One common answer is that they are pragmatically self-refuting: they are guaranteed to be false whenever held (Hintikka 1962; Williams 1994). For further discussion, see de Almeida (2001); Green and Williams (2007).

set of possibilities compatible with what the relevant agents believe. But why assume that there is a context-invariant answer to this question? In some contexts it may be the former, in others, the latter. And perhaps sometimes it's simply indeterminate.²⁰

While this seems reasonable, it doesn't help the classical semantics evade our challenge. After all, one candidate for the epistemic modal base is the set of possibilities compatible with what the relevant agents know. This predicts that there should be coherent readings of belief reports embedding epistemic contradictions. That is, there should be an available reading of (15) on which it's equivalent to (17), and an available reading of (16) on which it's equivalent to (18). Indeed, we'd expect listeners to converge on this reading, due to general principles of charity. However, as we have seen, there's no readily available coherent reading of belief reports embedding epistemic contradictions. Those who leave it to context to determine the epistemic state in question will have trouble explaining this observation.

5 Troubles for Non-Classical Semantics

For some readers, it may come as no surprise that the classical semantics has trouble explaining the infelicity of epistemic contradictions in belief reports. Yalcin (2007, 2011) has shown that epistemic contradictions are infelicitous when embedded under various operators, which poses a problem for the classical semantics. Compare:

- (21) # Suppose it's raining and it might not be raining.
- (22) ✓ Suppose it's raining and I don't know [/believe] it's raining.

One might hope for a unified treatment of epistemic contradictions in embedded contexts: whatever explains the infelicity of (21) also explains the infelicity of (1).

Indeed, a variety of non-classical semantics for *might* explain the infelicity of epistemic contradictions in embedded contexts. In doing so, they validate No Contradictions. However, we show that these theories are forced to deny Fallibility. For reasons of space, we focus on one implementation of a non-classical semantics for *might*: a version of the update semantics discussed in Veltman (1996); Gillies (2001); Yalcin (2012a,b); and Willer (2013). However, the problem we raise generalizes to other non-classical semantics, such as the static semantics developed by Yalcin (2007, 2011) and Moss (2015).

Update semantics is a type of dynamic semantics. In a dynamic semantics, the meaning of an expression is not its truth conditions. Rather, the meaning of an expression is its *context change potential*. This *ccp* is a function that takes as

²⁰This proposal would fit naturally with 'flexible contextualism' (Dowell 2011), according to which the community of relevant agents varies with the context of utterance.

input a context and returns the result of updating that context with the expression. In a slogan: the meaning of an expression is its ability to change a body of information.²¹

Consider a language \mathcal{L} containing a set of atomic sentences $\{\alpha_1, \dots, \alpha_n\}$ closed under *might* (\diamond), *and* (\wedge), and *not* (\neg). Let a possible world w be a function from atomic sentences to truth values. Let a context s be a set of possible worlds. According to update semantics, the interpretation of \mathcal{L} is a function $[\cdot]$ from sentences in \mathcal{L} to ccps, functions from contexts to contexts, defined recursively as follows:

Update Semantics

1. $s[\alpha] = s \cap \{w \mid w(\alpha) = 1\}$
2. $s[\phi \wedge \psi] = s[\phi][\psi]$
3. $s[\neg\phi] = s - s[\phi]$
4. $s[\diamond\phi] = \{w \in s \mid s[\phi] \neq \emptyset\}$.

According to this semantics, an atomic sentence narrows down a context to the worlds where it is true. A conjunction affects the context in two steps: first, it updates the context with the first conjunct; next, the resulting context is updated with the second conjunct. The negation of a sentence ϕ updates the context with the $\neg\phi$ worlds.

These first three clauses are updates: they affect the context by narrowing down the possible worlds in it. By contrast, *might* is a test. Rather than narrowing down the worlds in a context, $\diamond\phi$ checks whether the context is compatible with ϕ . For example, (13) (*The house might not be empty*) will leave the context unchanged provided there's at least one world in the context where the house isn't empty. Otherwise, the context crashes (represented as the empty set of worlds \emptyset).

This semantics predicts that unembedded epistemic contradictions are infelicitous (Veltman 1996; Gillies 2001). To see this, let us first define a notion of inconsistency for Update Semantics. A sentence ϕ is inconsistent just in case updating with ϕ is guaranteed to crash any context:

Consistency A sentence ϕ is consistent iff $\exists s: s[\phi] \neq \emptyset$; otherwise ϕ is inconsistent.

It is easy to show that epistemic contradictions are inconsistent in Update Semantics:

Fact 1 (Epistemic Contradictions are Inconsistent)

For any descriptive (non-modal) sentence ϕ and any context $s: s[\phi \wedge \diamond\neg\phi] = \emptyset$.

²¹For important contributions to the dynamic tradition, see Stalnaker (1973); Karttunen (1974); Heim (1982, 1983); Groenendijk and Stokhof (1990, 1991a).

Proof: Let s be an arbitrary context and ϕ an arbitrary descriptive sentence. By Update Semantics, $s[\phi \wedge \diamond\neg\phi] = s[\phi][\diamond\neg\phi]$. Now $s[\phi]$ is guaranteed to only contain ϕ worlds. Hence this set will always fail the test performed by $\diamond\neg\phi$. So $s[\phi \wedge \diamond\neg\phi] = \emptyset$.

To illustrate with an example, take (12): *The house is empty and it might not be empty*. Suppose our context s contains three worlds: w and u , in which the house is empty, and v , in which the house isn't empty. The first conjunct (*The house is empty*) narrows the context down to worlds where the house is empty, giving us $\{w, u\}$. The second conjunct then tests to see whether this updated context contains any worlds where the house isn't empty. Since the updated context fails this test, the sentence crashes. (See Figure 1.)




Figure 1: Updating with $\phi \wedge \diamond\neg\phi$

However, thus far Update Semantics doesn't make any predictions about our puzzle, because it lacks a semantics for *believes*. Let's now enrich \mathcal{L} with a *believes* operator (B_A). The standard dynamic semantics for *believes* analyzes belief in terms of *support*, where support is defined as a fixed point:

Support A context s supports ϕ ($s \models \phi$) iff $s[\phi] = s$.

On this definition, a context supports a sentence iff updating the context with that sentence has no effect on the context. According to the standard dynamic view, an agent believes ϕ iff her doxastic state supports ϕ .²² More precisely, suppose that an agent A 's doxastic state at a world w is characterized by a set of worlds (s_A^w): these are the worlds consistent with what A believes at w . Then:

Belief as Support $s[B_A\phi] = s \cap \{w \mid s_A^w \models \phi\}$.

To illustrate, consider:

(23) Ari believes the house is empty.

²²This account was first proposed by Hans Kamp, and is defended in Heim (1992); Zeevat (1992); Yalcin (2012a,b); Willer (2013).

According to Belief as Support, (23) narrows down a context to those worlds where Ari's doxastic state supports *the house is empty*, which in turn obtains iff there is no world in her doxastic state in which the house is occupied.

This approach has the advantage of validating No Contradictions. By Fact 1, epistemic contradictions are semantically inconsistent. And so no non-empty set of worlds supports an epistemic contradiction. *A fortiori*, no coherent agent's doxastic state supports an epistemic contradiction. By Belief as Support, it follows that no coherent agent believes an epistemic contradiction.²³

Unfortunately, Belief as Support invalidates Fallibility. To see this, consider:

(24) Ari believes the house might not be empty.

Given Update Semantics and Belief as Support, (24) updates the context with the information that Ari's doxastic state contains at least one world where the house is occupied. But this means that (23) is false. More generally, $B_A\phi$ and $B_A\Diamond\neg\phi$ provide incompatible instructions for updating the context. Thus Update Semantics, when combined with Belief as Support, forces us to abandon Fallibility.²⁴

Let's take stock. We've argued that the leading accounts of what it is to believe that something might be the case fail to satisfactorily resolve our puzzle. Specifically, we've canvassed two versions of a classical semantics for epistemic modals (the knowledge-based approach and the belief-based approach) as well as a standard dynamic approach (Update Semantics conjoined with Belief as Support). The knowledge-based approach failed to validate No Contradictions, whereas both the belief-based approach and the dynamic approach invalidated Fallibility.

Is there an alternative approach that can validate both Fallibility and No Contradictions? We think so. In what follows we present our own solution, which integrates Update Semantics with a Bayesian account of belief. We show

²³This strategy generalizes smoothly to explain the infelicity of epistemic contradictions under other attitude verbs. Take *supposes*. Update Semanticists can hold that *A supposes ϕ* narrows down the context to those worlds where *A*'s suppositional state supports ϕ . More precisely, let sup_A^w be the worlds compatible with what *A* supposes at w . Letting Su_A abbreviate *A supposes*, Update Semanticists can propose:

Supposes as Support $s[Su_A\phi] = s \cap \{w \mid sup_A^w \models \phi\}$.

Since no set of worlds supports an epistemic contradiction, epistemic contradictions cannot be supposed: (21) is always false. (Cf. the static treatment of *supposes* in Yalcin 2007.)

²⁴Other non-classical semantics for epistemic modals, such as those developed by Yalcin (2007, 2011) and Moss (2015), arrive at much the same impasse. For example, Yalcin (2007, 2011) relativizes the truth conditions of sentences to both a world w and an information state s . (As in Update Semantics, s is a set of worlds.) On Yalcin's semantics, $\Diamond\phi$ is true at some w, s iff s contains at least one world where ϕ is true. Yalcin combines this with a semantics for belief reports according to which *believes* shifts the information state to the believer's doxastic alternatives: $B_A\phi$ is true at some w, s iff for every world w' in s_A^w , ϕ is true at w', s_A^w . This predicts that (23) and (24) have incompatible truth conditions: (23) is true iff the house is empty at every world in Ari's doxastic state, and (24) is true iff the house is occupied at some world in Ari's doxastic state. So, much like Belief as Support, Yalcin's (2007, 2011) treatment of modal beliefs invalidates Fallibility.

that this position is able to validate all of our principles, thereby resolving the puzzle.

6 A New Semantics for Belief Reports

Here, in a nutshell, is our proposal: *an agent believes ϕ iff she assigns a sufficiently high credence to the result of updating her information with ϕ* . This proposal offers a way of synthesizing Update Semantics with a ‘Lockean’ approach to belief, according to which believing ϕ amounts to having a credence in ϕ that exceeds some threshold.²⁵ The resulting synthesis combines the primary advantages of both approaches, thereby resolving our puzzle. From Update Semantics, we borrow the resources to validate both No Contradictions and Uncertainty-Possibility Link. From the Lockean view, we borrow the resources to validate Uncertain Belief. And by validating both Uncertainty-Possibility Link and Uncertain Belief, we thereby validate Fallibility.

To introduce the details of our proposal, it may help to start with a simple version of the Lockean view, according to which A believes ϕ iff A assigns a sufficiently high credence to the set of ϕ worlds. Let $[\cdot]$ assign to each descriptive sentence of the language the set of worlds where it is true. Let Pr_A be A’s credence function, and let t denote some threshold between 0 and 1. Lockeans propose:

Lockean Belief $\llbracket B_A\phi \rrbracket^w = 1$ iff $Pr_A([\phi]) > t$.

That is, $B_A\phi$ is true iff A assigns a credence greater than t to the set of worlds where ϕ is true.

While we’ll argue shortly that Lockean Belief requires revision, the basic idea behind the Lockean approach holds considerable appeal. Unlike Belief as Support, Lockean Belief sheds light on the connection between outright belief and degrees of belief. In particular, it validates plausible inference patterns linking these two notions, for instance:

- (25) a. Fred believes it’s raining. \Rightarrow
 b. Fred is fairly [/quite] confident that it’s raining.

In addition, Lockean Belief validates Uncertain Belief. After all, it’s coherent to have a high credence in ϕ without being certain that ϕ . According to Lockean Belief, a high credence is all that’s required for belief.

Despite its attractions, Lockean Belief does not solve our puzzle. Taken by itself, it does not validate either Fallibility or No Contradictions. Indeed, taken by itself, it is not clear how it accounts for beliefs about what might be the case. This lacuna is particularly evident if we adopt the dynamic approach to epistemic modals from §5. After all, on the dynamic approach there is no set of worlds where $\Diamond\phi$ is true. Thus the task that now faces us is to extend Lockean Belief to modal beliefs in a way that validates our two principles.

²⁵For defenses of the Lockean view, see [Foley \(1993\)](#); [Christensen \(2005\)](#); [Sturgeon \(2008\)](#).

To do this, we propose a dynamic twist to Lockean Belief. On a dynamic approach, while there is no set of worlds in which $\diamond\phi$ is true, updating any particular set of worlds s with $\diamond\phi$ will always result a set of worlds (either s or \emptyset). Thus to capture modal beliefs, we propose an updated Lockean thesis: for A to believe ϕ is to assign a high credence to the result of updating A 's doxastic state with ϕ . Intuitively, we can think of updating with one's doxastic state with ϕ as a way of modeling becoming certain of ϕ . Given this gloss, our proposal amounts to the following. A believes ϕ iff A assigns a sufficiently high credence to the doxastic state that would result from becoming certain of ϕ .

In order to implement this, we model an agent A 's doxastic state at w with two components—a set of worlds s_A^w , and a probability function Pr_A^w . s_A^w is the set of worlds compatible with what A is certain of at w . As in Lockean Belief, Pr_A^w is an agent's credence function at w . It assigns s_A^w a probability of 1. We model the result of updating A 's information at w with ϕ as $s_A^w[\phi]$, where $[\cdot]$ is defined as in Update Semantics. We propose that an agent believes ϕ iff she assigns a sufficiently high credence to *this* set of worlds ($s_A^w[\phi]$). And so updating a context with $B_A\phi$ narrows down that context to the worlds where A 's credence function meets this condition. More precisely:

Locke Updated $s[B_A\phi] = s \cap \{w \mid Pr_A^w(s_A^w[\phi]) > t\}$.

Let's unpack this. According to Locke Updated, $B_A\phi$ updates a context with the worlds where A believes ϕ . Which worlds are these? The worlds where A has a sufficiently high credence in $s_A^w[\phi]$. Here $s_A^w[\phi]$ represents the doxastic state that A would be in, if A were to become certain of ϕ .²⁶

We now show how this semantics resolves our puzzle. To do this, we first show that our semantics validates both Uncertain Belief and Uncertainty-Possibility Link, thereby validating Fallibility. We next show that our semantics validates No Contradictions.

Start with Uncertain Belief. To see that Locke Updated validates this principle, note that Locke Updated agrees with Lockean Belief when it comes to descriptive beliefs. If ϕ is descriptive, then to believe ϕ is to assign a sufficiently high credence to the ϕ worlds:

Fact 2 (Descriptive Beliefs Are Lockean)

For any agent A and any descriptive sentence ϕ : $s[B_A\phi] = s \cap \{w \mid Pr_A^w([\phi]) > t\}$.

Proof: By Locke Updated, $B_A\phi$ holds at a world w iff A 's credence in $s_A^w[\phi]$ exceeds t . To find $s_A^w[\phi]$, we take the set of worlds in A 's doxastic state at w (s_A^w) and update this set with ϕ . By Update Semantics, when ϕ is descriptive, this

²⁶Locke Updated can be complicated in various ways. For example, one could also allow the threshold to vary with the context of utterance and/or the believer's practical interests (Weatherston 2005, 2012; Ganson 2008; Fantl and McGrath 2009). One could also make *believes* sensitive to a question under consideration (Yalcin 2011, forthcoming). For our purposes, we set these complications aside, since they are not directly relevant to our puzzle.

is simply the result of intersecting s_A^w with the ϕ worlds ($s_A^w \cap \llbracket \phi \rrbracket$). Since every agent assigns credence 1 to the set of worlds in her doxastic state, her credence in $\llbracket \phi \rrbracket$ will equal her credence in $s_A^w[\phi]$.

From Fact 2, it's a short step to Uncertain Belief. On a simple and quite natural view, an agent is certain of ϕ just in case her doxastic state supports ϕ . This in turn entails that her credence in ϕ is 1. Let C_A be an operator short for *A is certain that*. This gives us the following:

Certainty as Support $s[C_A\phi] = s \cap \{w \mid s_A^w \models \phi\}$.²⁷

Thus, on Locke Updated, anyone whose credence in a descriptive claim ϕ is greater than t but less than 1 will count as believing ϕ without being certain of it.

To apply our semantics, recall Ari, who believes the house is empty without being certain of it. We model this by saying that, at every world in the context, Ari's doxastic state includes both worlds where the house is empty and worlds where it isn't. To simplify, suppose Ari's doxastic state consists of just our three worlds from §5: w and u , in which the house is empty, and v , in which there's someone inside. Suppose that Ari assigns a credence of .8 to $\{w, u\}$ and a credence of .2 to $\{v\}$. Finally, suppose that t , the threshold for belief, is .75. Given all of this, Locke Updated entails that (23) (repeated here as (26)) is supported:

(26) Ari believes the house is empty.

After all, the result of updating Ari's doxastic state with *the house is empty* is $\{w, u\}$, and Ari's credence in this proposition exceeds .75.

At the same time, our semantics for *certain* predicts that the following will update the context to return the empty set:

(27) Ari is certain that the house is empty.

After all, Ari's doxastic state includes a world where the house is not empty (v).

Let's turn to Uncertainty-Possibility Link. To see that our semantics validates this principle, note that while Locke Updated agrees with Lockean Belief when it comes to descriptive beliefs, they diverge when it comes to beliefs about what might be the case. In particular, Locke Updated agrees with Belief as Support that modal beliefs are 'transparent': an agent believes $\diamond\phi$ just in case her doxastic state contains a ϕ world. Summarizing:

Fact 3 (Might Beliefs Are Transparent)

For any agent A and any descriptive sentence ϕ : $s[B_A\diamond\phi] = s \cap \{w \mid s_A^w \not\models \neg\phi\}$.

²⁷In §3.1, we suggested that the truth-value of a certainty ascription depends on the context (either that of the conversation, or that of the subject of the certainty ascription). One way to implement this suggestion in the present framework would be to let the subject's doxastic alternatives vary with context, so that in high-stakes contexts s_A^w includes possibilities that are absent in low-stakes contexts. See Clarke (2013) for an account of belief along much these lines.

Proof: By Locke Updated, A believes $\diamond\phi$ at w just in case she gives sufficiently high credence to $s_A^w[\diamond\phi]$. By Update Semantics, $s_A^w[\diamond\phi]$ is either s_A^w or \emptyset , depending on whether there is a ϕ world in s_A^w . If there is, then $s_A^w[\diamond\phi] = s_A^w$, to which A assigns credence 1. Otherwise, $s_A^w[\diamond\phi] = \emptyset$, to which A assigns credence 0. And so A believes $\diamond\phi$ just in case her doxastic state includes a ϕ world.

Fact 3 and Certainty as Support entail Uncertainty-Possibility Link. Given Certainty as Support, if A isn't certain that ϕ , then A's doxastic state doesn't support $\neg\phi$. And so, given Fact 3, A believes $\diamond\phi$. To illustrate, let's return to Ari. Since Ari isn't certain that the house is empty, her doxastic state contains a world where the house isn't empty ($\{v\}$). From Fact 3, we derive that Ari believes the house might not be empty.

Because our semantics validates Uncertain Belief and Uncertainty-Possibility Link, it validates Fallibility as an immediate corollary. According to our semantics, anytime an agent believes ϕ without being certain that ϕ , she will also count as believing $\diamond\neg\phi$. We summarize Facts 2 and 3 in Figure 2, which illustrates the different constraints imposed by being certain that ϕ , believing ϕ , and believing $\diamond\phi$.




Figure 2: Locke Updated

Finally, our semantics validates No Contradictions:

Fact 4 (No Contradictions)

For any context s , agent A and any descriptive sentence ϕ : $s[B_A(\phi \wedge \diamond\neg\phi)] = \emptyset$.

Proof: By Locke Updated, A believes $(\phi \wedge \diamond\neg\phi)$ at w iff A assigns a sufficiently high credence to $s_A^w[\phi \wedge \diamond\neg\phi]$. From Fact 1, we know that for any context s , $s[\phi \wedge \diamond\neg\phi] = \emptyset$. So $s_A^w[\phi \wedge \diamond\neg\phi] = \emptyset$. Consequently, $Pr_A^w(s_A^w[\phi \wedge \diamond\neg\phi]) = 0$.

Applied to our example, (1) (*Ari believes the house is empty and might not be empty*) is true iff Ari assigns a sufficiently high credence to the result of updating her doxastic state with *the house is empty and might not be empty*. This proposition is found by taking her doxastic state ($\{w, u, v\}$) and updating it in two steps. First, we update it with *the house is empty*, giving us $\{w, u\}$. Next,

we update this set with *the house might not be empty*, which requires checking whether $\{w, u\}$ contains at least one world where the house isn't empty. Since there's no such world in $\{w, u\}$, we get the empty set. Since coherent agents assign credence 0 to the empty set, Ari cannot coherently believe that the house is empty and might not be.

7 Against Closure

Our solution to the puzzle has a surprising consequence: rational belief is not closed under logical implication.

Here's an off-the-shelf formulation of a multi-premise closure principle:

Multi-Premise Closure (MPC) If (i) A is rational in believing premises $\phi_1 \dots \phi_n$, (ii) $\phi_1 \dots \phi_n \models \psi$, and (iii) A competently infers ψ from these premises, then A's resulting belief in ψ is rational.²⁸

Let ϕ_1 be the premise: *the house is empty*. Let ϕ_2 be the premise: *the house might not be empty*. Let ψ be the conclusion: *the house is empty and it might not be empty*. On our account, Ari can coherently (and presumably, rationally) believe ϕ_1 and ϕ_2 , but it would be incoherent (and hence irrational) for her to believe ψ on this basis. Indeed, this is not just an idiosyncratic feature of our semantics: any semantics that validates both Fallibility and No Contradictions is forced to reject MPC.²⁹

Of course, challenges to closure are nothing new: it's well-known that lotteries (Kyburg 1961) and prefaces (Makinson 1965) cause trouble for MPC. But if we're right, counterexamples to MPC are much more common than has been acknowledged. What our account suggests is that anytime you believe something without being certain of it, there's a counterexample to MPC lurking. If—as seems plausible—relatively few of our beliefs count as certain, this means that most of our beliefs furnish us with counterexamples to MPC.

Here is another way in which our counterexample to MPC is stronger than others. While Bayesian theories of belief reject MPC, they accept a weaker principle (Adams 1966; Edgington 1997; Sturgeon 2008). Say that an agent's uncertainty in ϕ is the difference between 1 and her credence in ϕ . So if an agent is certain that ϕ , her uncertainty in ϕ is 0. And if an agent is certain that $\neg\phi$, then her uncertainty in ϕ is 1. Bayesians accept:

Bayesian Closure (BC) If $\phi_1 \dots \phi_n \models \psi$, then a rational agent's uncertainty in ψ cannot be greater than the sum of her uncertainty with respect to ϕ_1, ϕ_2, \dots , and ϕ_n .

Our example is also a counterexample to BC. Since Ari's credence in ϕ_1 (*the house is empty*) is .8, her uncertainty with respect to ϕ_1 is .2. So Ari should be

²⁸Our formulation is based on Schechter (2013).

²⁹Here, we make the assumption that $\phi; \diamond\neg\phi \models \phi \wedge \diamond\neg\phi$. This holds on standard dynamic notions of validity such as 'update-to-test' and 'test-to-test' entailment (Veltman 1996).

certain in ϕ_2 (*the house might not be empty*), which means her degree of uncertainty with respect to ϕ_2 should be 0. However, Ari should also be certain that ψ (*the house is empty and it might not be empty*) is false. So Ari’s degree of uncertainty in ψ is 1, which exceeds the sum of her uncertainty in the premises ϕ_1 (.2) and ϕ_2 (0).³⁰

For those attracted to closure, it’s natural to try to restrict MPC so that it does not apply to reasoning involving epistemic modals:

Restricted MPC If (i) $\phi_1 \dots \phi_n$ and ψ are descriptive, (ii) A is rational in believing $\phi_1 \dots \phi_n$, (iii) $\phi_1 \dots \phi_n \models \psi$, and (iv) A competently infers ψ from these premises, then A’s resulting belief in ψ is rational.

Restricted MPC captures many of the intuitions that motivated closure in the first place—in particular, the idea that deduction is a rational way of extending our beliefs. On the picture that emerges, deduction is always a rational way of extending our beliefs about what *is* the case, but it’s not always a rational way of extending our beliefs about what *might* be the case.

Of course, validating Restricted MPC without invalidating Uncertain Belief is no easy task. Locke Updated is not up to it. This is because Locke Updated, like Lockean Belief, embraces the idea that belief only requires meeting some threshold less than 1. Notoriously, any such Lockean approach stands in tension with even Restricted MPC.³¹

However, there are promising strategies for modifying the Lockean thesis to preserve closure. For example, Leitgeb defends a ‘stability’ theory of belief, according to which “Belief is determined by a proposition of resiliently or stably high subjective probability” (2014: 145). The intuitive idea here is that in order

³⁰Our counterexample to MPC has affinities with a recent counterexample to *single-premise* closure developed independently by Bledin and Lando (forthcoming). Bledin and Lando’s counterexample relies on what they (following Yalcin 2007) refer to as “Łukasiewicz’s Principle”:

Łukasiewicz’s Principle $\neg\phi \models \neg\Diamond\phi$

Proponents of non-classical semantics for epistemic modals often embrace a non-classical consequence relation that validates Łukasiewicz’s Principle (Veltman 1996; Yalcin 2007). But, as Bledin and Lando observe, it’s not rational to believe $\neg\Diamond\phi$ on the basis of a belief in $\neg\phi$. For example, Ari rationally believes the house is empty. But since she isn’t certain of this, it isn’t rational for her to believe that the house *must* be empty.

While both of our counterexamples involve epistemic modals, it is worth highlighting an important difference. Unlike Bledin and Lando’s counterexample, our counterexample does not require embracing any non-classical consequence relation, let alone one that validates Łukasiewicz’s Principle. All our counterexample requires is that conjunction introduction is a valid form of argument. This difference is important, because it creates difficulty for a potential response to Bledin and Lando’s counterexample: namely, to perform *modus tollens*, and reject any semantics that validates Łukasiewicz’s Principle (Schulz 2010: 388). Our counterexample shows that in order to preserve closure one would also have to reject conjunction introduction. Since conjunction introduction is classically valid, we suspect that many would find such a rejection to be much more unpalatable.

Thus while our counterexample is strictly independent of Bledin and Lando’s, we take our counterexample to complement theirs. Indeed, Locke Updated can explain why it’s not always rational to reason in accordance with Łukasiewicz’s Principle: it’s only rational to believe $\neg\Diamond\phi$ if one is certain that $\neg\phi$ is true.

³¹However, Locke Updated does validate a restricted form of BC, according to which BC holds for descriptive premises and conclusions.

to count as believing ϕ , an agent's credence in ϕ must remain sufficiently high even upon acquiring new information (within certain limits). Leitgeb develops this idea in terms of a technical notion of 'P-Stability'. Say that a sentence ϕ is *P-stable* relative to a probability function Pr iff $Pr(\llbracket\phi\rrbracket)$ remains higher than $\frac{1}{2}$ even after conditionalizing on any *admissible* claim. And say that ψ is *admissible* (relative to ϕ , Pr) iff ψ is consistent with ϕ and $\llbracket\psi\rrbracket$ is assigned some non-zero probability by Pr . That is, the set of P-stable sentences for Pr ($\mathcal{P}(Pr)$) can be characterized as follows:

P-stability $\mathcal{P}(Pr) = \{\phi \mid \forall \psi \in \mathcal{A}(\phi, Pr): Pr(\llbracket\phi\rrbracket \mid \llbracket\psi\rrbracket) > \frac{1}{2}\},$
 where $\psi \in \mathcal{A}(\phi, Pr)$ iff $\phi; \psi \not\vdash \perp$, and $Pr(\llbracket\psi\rrbracket) > 0$.

(When ϕ is descriptive, this is equivalent to saying that ϕ is P-stable just in case every world at which ϕ is true is assigned a higher probability than the union of the worlds at which ϕ is false.)

As Leitgeb develops the stability theory, an agent believes ϕ just in case ϕ is entailed by some P-stable claim ψ that she believes, and hence her credence in ϕ is greater than or equal to her credence in ψ .³² Leitgeb shows that, on such a view, an agent's beliefs will be closed under logical implication.

As stated, Leitgeb's view is schematic. After all, a given credence function can generate multiple P-stable sentences. And for every such P-stable sentence ϕ , an agent's credence in $\llbracket\phi\rrbracket$ gives us a different candidate for the Lockean threshold. To convert Leitgeb's view into a semantics for *believes*, we can posit a contextually determined choice function f that, given the agent's set of P-stable propositions, selects one as the Lockean threshold.³³ Those attracted to this approach could modify Locke Updated as follows:

Locke Stabilized $s[B_A\phi] = s \cap \{w: Pr_A^w(s_A^w[\phi]) \geq Pr_A^w(f(\mathcal{P}(Pr_A^w)))\},$
 where f selects a unique member of $\mathcal{P}(Pr_A^w)$.

To illustrate, consider (23) (*Ari believes the house is empty*). According to Locke Stabilized, this means that Ari's credence in the result of updating her doxastic state with the claim that the house is empty exceeds that of the f -selected P-stable proposition.

While this semantics validates Restricted MPC, it does not validate an unrestricted closure principle: MPC will still fail when it comes to ϕ and $\diamond\neg\phi$. After all, on Locke Stabilized (as on Locke Updated) an agent believes $\diamond\phi$ iff ϕ is compatible with her doxastic state. In that case, $s_A^w[\diamond\phi] = s_A^w$. And since $Pr_A^w(s_A^w) = 1$, $Pr_A^w(s_A^w[\llbracket\psi\rrbracket]) = 1$ for every ψ . But $\phi \wedge \diamond\neg\phi$ still crashes A 's doxastic state, and so she does not believe this.

For our purposes, we need not commit to a stability theory of belief. Perhaps the reader prefers some other closure-preserving modification of a Lockean view. If so, we should be able to import any such modification into our

³²Technically, Leitgeb takes P-stability to be a property of propositions. However, since we are interested in giving a semantics for belief reports that encompasses modal beliefs (which are not themselves characterized via propositions), we need to treat P-stability as a property of sentences.

³³See Leitgeb (2013) for the view that f selects the strongest P-stable proposition.

semantics for *believes*, thereby preserving Restricted MPC. Or perhaps the right response to lotteries and paradoxes is to abandon even restricted closure principles. For our purposes, the important point is that beliefs involving epistemic modals provide strong grounds for abandoning (unrestricted) MPC; we leave it as an open question whether our semantics for *believes* should validate Restricted MPC (and, if so, how this validation is best achieved).

8 Objections

8.1 First Objection: Order Effects

While Update Semantics predicts the semantic inconsistency of epistemic contradictions, it does not predict the semantic inconsistency of *reversed epistemic contradictions*, i.e., sentences of the form:

$$(28) \quad \Diamond \neg \phi \wedge \phi$$

For example:

$$(29) \quad ? \text{ The house might not be empty and it is empty.}$$

To see this, recall our context s containing three worlds— w and u , in which the house is empty, and v , in which it isn't empty. Updating with the first conjunct of (29) (*The house might not be empty*) leaves the context the same, while updating with the second conjunct shrinks it down to $\{w, u\}$. So $s[(29)] \neq \emptyset$.

As a result, our solution to the puzzle does not predict that it's incoherent to believe reversed epistemic contradictions. That is, our semantics does not predict that instances of the following are infelicitous:

$$(30) \quad B_A(\Diamond \neg \phi \wedge \phi)$$

To see this, suppose (as before) that Ari's doxastic state is $\{w, u, v\}$. Then:

$$(31) \quad ? \text{ Ari believes that the house might not be empty and (it) is empty.}$$

is predicted to narrow down the context to those worlds in which Ari's credence in $\{w, u\}$ exceeds the Lockean threshold.

Some may regard this as a problem for our approach. After all, (31) sounds fairly odd. Thus some might insist that any adequate solution to our puzzle will explain not just the incoherence of believing epistemic contradictions, but also the incoherence of believing reversed epistemic contradictions.

In response, we should first note that there is a delicate and—to our knowledge—currently unresolved question as to what exactly the data are. Historically, many dynamic semanticists have regarded it as a datum that discourse (32) is coherent, or at least less degraded than (33):

$$(32) \quad \text{It might be raining. It isn't raining.}$$

(33) ? It isn't raining. It might be raining.³⁴

In a similar vein, [Sorensen \(2009\)](#) and [Dorr and Hawthorne \(2013\)](#) claim that reversing the order of the conjuncts of an embedded epistemic contradiction tends to make the sentence more acceptable. Others have questioned these judgments.³⁵ Until this empirical issue is investigated more fully, it's not obvious whether the fact that our theory predicts that (30) is consistent should be regarded as a vice or a virtue.

But suppose we set aside this question about the data and assume, at least for the sake of argument, that reversed epistemic contradictions are typically judged infelicitous. Is there any way of modifying Update Semantics to predict this? Note that while reversed epistemic contradictions are semantically consistent, they do display a somewhat unusual property: they are *non-idempotent*, where idempotence is defined as follows:

Idempotence A sentence ϕ is idempotent iff for any context s , $s[\phi] \models \phi$.

To put it in intuitive terms: a sentence is idempotent just in case updating any context with the sentence once achieves the same result as updating the context with the sentence twice. To see that reversed epistemic contradictions are non-idempotent, consider again our context s containing two ϕ worlds (w and u) and a $\neg\phi$ world (v). As we've seen, $s[\diamond\neg\phi \wedge \phi] = \{w, u\}$; however, $s[\diamond\neg\phi \wedge \phi][\diamond\neg\phi \wedge \phi] = \emptyset$. As [Yalcin \(2015\)](#) remarks, there seems to be something odd about a sentence that behaves in this way: uttering it creates a context that is "inhospitable for its own update" (502). Thus one option is to modify Update Semantics to predict that non-idempotent sentences always crash the context.³⁶

We do not wish to take a stand on whether Update Semantics should be modified in this way. For our purposes, it suffices to observe that the present objection to our proposal raises larger issues about whether Update Semantics should be modified in order to predict that reversed epistemic contradictions are infelicitous. For those who think that some such modification is required, a natural path is to locate the infelicity of reversed epistemic contradictions in their non-idempotence.

8.2 Second Objection: Questioning Fallibility

According to Locke Updated, believing a conjunction is not equivalent to believing each conjunct individually. In particular, while Locke Updated predicts that $B_A(\phi \wedge \diamond\neg\phi)$ crashes, it doesn't predict that (34) crashes:

³⁴See, for example, [Groenendijk and Stokhof \(1991b\)](#); [Veltman \(1996\)](#); [Gillies \(2001\)](#); [von Fintel and Gillies \(2007\)](#).

³⁵See, for example, [Yalcin \(2015\)](#).

³⁶It is a delicate matter how best to formulate this modification. As Yalcin notes, some quantified non-idempotent sentences seem to be felicitous, which suggests that a blanket ban on all non-idempotent sentences is too strong. Drawing on [Klinedinst and Rothschild \(2015\)](#), Yalcin offers a weaker non-idempotence requirement ('Consecutive update idempotence') that applies at the compositional semantic level. This weaker requirement suffices to explain the infelicity of (31). (See [Yalcin 2015](#): 503-504 for the details.)

(34) $B_A\phi. B_A\Diamond\neg\phi$.

Many find discourses like this at least somewhat odd. Consider, for instance:

(35) ? Ari believes the house is empty. She also believes it might not be empty.

Should we revise our semantics to make (35) inconsistent after all?

Here too there is a question about the data. While we agree that (35) is a bit peculiar, it doesn't sound to us *as* bad as (1). Informal polling suggests others agree: roughly half of those we've surveyed judge (35) to be marked, while the other half judged it to be fine. By contrast, the vast majority of respondents deem (1) infelicitous. This provides reason to resist revising our semantics to predict that instances of (34) crash: an adequate solution to the puzzle will capture the fact that (1) sounds worse than (35).

In addition to our scruples about the infelicity of (35), there are principled reasons to resist revising our semantics to predict that instances of (34) crash. Any such revision would require giving up Fallibility. But, as we have seen (§3.1), there are three independent arguments for Fallibility. To review: there was the argument from the felicity of CBAs (e.g., (3)); there was the argument from the felicity of variants of (i) involving *realize* and *recognize* (e.g., (4)); finally, there was the theoretical argument, which derived Fallibility from the conjunction of Uncertain Belief and Uncertainty-Possibility Link, both of which were supported by independent data.

For ease of reference, (36) collects much of the relevant data in one place, showcasing these three grades of modal infelicity:

- (36) a. # A believes $(\phi \wedge \Diamond\neg\phi)$.
 b. ? A believes ϕ . A also believes $\Diamond\neg\phi$.
 c. ✓ A believes ϕ . But A realizes $\Diamond\neg\phi$.
 d. ✓ I believe ϕ . But $\Diamond\neg\phi$.

None of the views that we've considered predicts all of these judgments. However, Locke Updated comes the closest. On the one hand, views that reject Fallibility (e.g., the belief-based version of the classical semantics and Belief as Support) incorrectly predict that all of these sentences are incoherent. On the other hand, the knowledge-based version of the classical semantics doesn't predict that any of these sentences are incoherent. Only Locke Updated predicts both the infelicity of (36a) and the felicity of (36c)-(36d).

Still, our initial question remains: what explains why (35/36b) is at least somewhat odd? We are not sure, but one hypothesis is that optional modal subordination is responsible.³⁷ A modal is subordinated when it is evaluated relative to a range of possibilities controlled by previous discourse. We propose that when we evaluate (35/36b), we tend to access a reading that subordinates

³⁷For discussion of modal subordination, see Roberts (1989); Kibble (1995); van Rooij (2005). For evidence that *might* gives rise to optional modal subordination, see Klecha (2012).

the modal *might* to the complement of the previous belief report (*the house is empty*). On this reading, the modal is evaluated relative to a set of worlds where the house is empty; consequently, the discourse crashes. While this reading is available, it is not mandatory; there’s another reading of (35/36b) on which the modal is not subordinated. On this reading, the discourse does not crash. We suspect the availability of this unsubordinated reading explains why (35) is less odd than (1).

This explanation leaves some questions unanswered. First, why do some speakers prefer the subordinated reading over the unsubordinated reading? Second, why does (4/36c) sound better than (35/36b)? We do not at present have fully developed answers to these questions. We suspect that when a modal has two readings—one subordinated, the other unsubordinated—the extent to which one reading will be preferred over the other will be influenced by a variety of factors. Perhaps for some speakers subordinated readings are the default—the readings they tend to latch onto in the absence of cues to the contrary. But this default can be overridden. Tellingly, (4/36c) contains a contrast marker (*but*); what’s more, this discourse sounds best when the contrast marker is given prosodic focus. We suspect that this may override the default, biasing speakers towards the unsubordinated reading.³⁸ However, we will leave to future research the project of developing a rigorous account of the factors that influence the extent to which the different readings are preferred. What’s important for our purposes is that there are compelling reasons to retain Fallibility. This suggests that the right response to (35/36b) is not to revise Locke Updated, but rather to explain the oddity of these sentences via some pragmatic mechanism.

8.3 Third Objection: *Might* Beliefs vs. *Might* Certainties

According to our view, an agent believes $\diamond\phi$ iff ϕ is compatible with her certainties. Given Certainty as Support, an agent is certain that $\diamond\phi$ under the exact same conditions. Thus our proposal collapses believing ϕ is possible and being certain that ϕ is possible:

Fact 5 (Collapse)

For any context s , agent A and any descriptive sentence ϕ : $s[B_A \diamond\phi] = s[C_A \diamond\phi]$.

Proof: From Fact 2, $s[B_A \diamond\phi] = s \cap \{w \mid s_A^w \not\models \neg\phi\}$. By Certainty as Support and Update Semantics, $s[C_A \diamond\phi] = s \cap \{w \mid s_A^w \not\models \neg\phi\}$. So $s[B_A \diamond\phi] = s[C_A \diamond\phi]$.

This seems counterintuitive. Consider, for example, DeRose’s (1991) cancer case, in which Jane’s husband John has undergone a test to determine whether he has cancer. A negative result will mean that John definitely does not have cancer. A positive result does not necessarily mean that John does have cancer;

³⁸It’s well known that contrast—and coherence relations more generally—influence how we resolve ambiguous sentences (Hobbs 1985; Kehler 2002). For discussion of how coherence relations influence the resolution of epistemic modals in particular, see Asher and McCready (2007).

rather, it means that further tests have to be run. It seems natural to describe Jane's credal state as follows:

- (37) Jane believes John might have cancer. But she isn't certain he might have cancer.

One possible response is to claim that the modal in the second conjunct is not epistemic; rather, it quantifies over physical or metaphysical possibilities (Stephenson 2007b: 50). However, some might regard this as *ad hoc*. At the very least, it doesn't seem obvious that the occurrence of *might* in the second conjunct is non-epistemic.

Perhaps a better response is to concede the counterexample and amend our semantics for certainty ascriptions. Earlier we raised the possibility that there is a stability constraint on belief. We could likewise impose a stability constraint on certainty. According to this constraint, in order for A to be certain that ϕ , A's doxastic state must support ϕ even once it's been updated with any admissible claim. As before, ψ is admissible iff ψ is consistent with ϕ and assigned some non-zero credence by A. That is:

Certainty Stabilized $s[C_A\phi] = s \cap \{w \mid \forall \psi \in \mathcal{A}(\phi, Pr_A^w): s_A^w[\psi] \models \phi\}$,

where $\psi \in \mathcal{A}(\phi, Pr_A^w)$ iff $\phi; \psi \not\models \perp$ and $Pr_A^w(\llbracket \psi \rrbracket) > 0$.³⁹

To see how this solves the present difficulty, let ψ be the claim that the test results are negative. This is consistent with the claim that John might have cancer; Jane also assigns this claim some non-zero credence. But if we update Jane's doxastic state with ψ , the resulting set contains no worlds in which John has cancer. And so the resulting set will fail the test imposed by the sentence: *John might have cancer*. Thus Certainty Stabilized predicts that the second sentence in (37) (*She isn't certain he might have cancer*) is true, as desired.

Certainty Stabilized is a conservative extension of Certainty as Support. That is, the two semantics make the same predictions whenever ϕ is not a possibility claim. To see why, we need to introduce some further terminology. A sentence ϕ is *persistent* when any context s that supports ϕ will continue to support ϕ once s is updated with more information (Veltman 1996: 3):

Persistence ϕ is persistent iff $\forall s \forall \psi: \text{if } s \models \phi, \text{ then } s[\psi] \models \phi$.

In Update Semantics, any descriptive claim is persistent. After all, s supports some descriptive claim ϕ just in case ϕ holds at every world in s . Whenever this obtains, ϕ also holds at any subset of s . By contrast, $\diamond\phi$ is not persistent (Veltman 1996). After all, $s \models \diamond\phi$ as long as updating s with ϕ doesn't produce \emptyset . It is consistent with this that updating certain subsets of s with ϕ will produce \emptyset . For example, when $s \models \diamond\phi$ we will not in general have that $s[\neg\phi] \models \diamond\phi$.

³⁹This corresponds to the claim that A is certain of ϕ iff ϕ is *P¹-stable*, where ϕ is *P¹-stable* iff, for any sentence ψ that is both consistent with ϕ and assigned some non-zero by A, $Pr(s_A^w[\phi] \mid \llbracket \psi \rrbracket) = 1$. For further discussion of different levels of P-stability, see Leitgeb (forthcoming: appendix B).

It turns out that whenever ϕ is persistent, Certainty Stabilized agrees with Certainty as Support. That is, assuming the semantics for $C_A\phi$ is provided by Certainty Stabilized, the following holds:

Fact 6

For any agent A and any persistent sentence ϕ :

$$s[C_A\phi] = s \cap \{w \mid s_A^w \models \phi\}.$$

Proof: It suffices to show that if ϕ is persistent, then $s_A^w \models \phi$ iff $\forall \psi \in \mathcal{A}(\phi, Pr_A^w) : s_A^w[\psi] \models \phi$. So suppose $s_A^w \models \phi$. Then since ϕ is persistent, $\forall \psi \in \mathcal{A}(\phi, Pr_A^w) : s_A^w[\psi] \models \phi$. Similarly, suppose that $\forall \psi \in \mathcal{A}(\phi, Pr_A^w) : s_A^w[\psi] \models \phi$. Then since $\top \in \mathcal{A}_\phi, s_A^w \models \phi$.

This last fact leads to an interesting prediction. While $\diamond\phi$ is not persistent in Update Semantics, $\Box\phi$ is.⁴⁰ After all, $s \models \Box\phi$ just in case $s \models \phi$. If this holds, then by the persistence of ϕ we know that for any $\psi : s[\psi] \models \phi$, and hence that $s[\psi] \models \Box\phi$. We've seen that while $\diamond\phi$ is not persistent, $\Box\phi$ is. Thus while Certainty Stabilized allows for a distinction between belief and certainty about what might be the case, it doesn't allow for any distinction between belief and certainty about what *must* be the case. Interestingly, this prediction appears to be born out by the data. Consider the variant of (37) that replaces *might* with *must*:

- (38) ?? Jane believes John must have cancer. But she isn't certain he must have cancer.

We find it much harder to access a true reading of (38) than (37). This is a surprising observation—one that Certainty Stabilized elegantly explains.

9 Conclusion

Recent work on the semantics of epistemic modals has explored what sort of mental state is involved in beliefs about epistemic possibility: what is it to believe that something *might* be the case? In this paper, we've tried to make progress on this question. We began by identifying two principles that should constrain any account of what's involved in believing that something might be the case. Taken together, these principles form a puzzle, since they are—at first blush—difficult to reconcile. We went on to resolve this puzzle by offering a new semantics for *believes* that integrates a Bayesian approach to belief with a dynamic semantics for modals.

While we have focused on belief, our approach extends to other attitudes. Say that an attitude verb V is *credal* iff $V_A\phi$ entails that A assigns some credence to ϕ . Many—perhaps most—attitude verbs are credal: one cannot suspect, regret, fear, or hope the house is empty if one is certain the house isn't empty. A natural generalization of our semantics for *believes* holds that for any credal

⁴⁰Here we assume that \Box is the dual of \diamond , and so $s[\Box\phi] = \{w \in s \mid s \models \phi\}$.

attitude verb V , $V_A\phi$ entails that A assigns some credence to the result of updating A 's doxastic state with ϕ (i.e., $Pr_A^w(s_A^w[\phi]) > 0$).

This proposal provides a general explanation of the oddity of epistemic contradictions under credal attitude verbs:

- (39) # Ari suspects/regrets/fears/hopes that [the house is empty and might not be].

After all, updating A 's doxastic state with an epistemic contradiction always returns the empty set. Since A will—if coherent—always have credence 0 in the empty set, this explains the incoherence of adopting any credal attitude towards an epistemic contradiction.⁴¹ Thus the account of belief developed in this paper has repercussions for our understanding of a broader class of attitudes.

References

- Ernest Adams. Probability and the logic of conditionals. In Hintikka and Suppes, editors, *Aspects of Inductive Logic*, pages 165–316. North-Holland, Amsterdam, 1966.
- Pranav Anand and Valentine Hacquard. Epistemics and attitudes. *Semantics and Pragmatics*, 6:1–59, 2013.
- D.M. Armstrong. *Belief, Truth, and Knowledge*. Cambridge University Press, London, 1973.
- Nicholas Asher and Eric McCready. Were, would, might and a compositional account of counterfactuals. *Journal of Semantics*, 24:93–129, 2007.
- A.J. Ayer. *Language, Truth, and Logic*. Dover, New York, 1936.
- Alexander Bird. Justified judging. *Philosophy and Phenomenological Research*, 74:81–110, 2007.
- Justin Bledin and Tamar Lando. Closure and epistemic modals, forthcoming.
- David Christensen. *Putting Logic in its Place*. Oxford University Press, 2005.
- Roger Clarke. Belief is credence one (in context). *Philosophers' Imprint*, 13(11):1–18, 2013.
- Claudio de Almeida. What moore's paradox is about. *Philosophy and Phenomenological Research*, 62(1):33–58, 2001.
- Keith DeRose. Epistemic possibility. *Philosophical Review*, 100:581–605, 1991.
- Dylan Dodd. Against fallibilism. *Australasian Journal of Philosophy*, 89(4):665–685, 2011.
- Dylan Dodd. Belief and certainty. *Synthese*, forthcoming.
- Cian Dorr and John Hawthorne. Embedding epistemic modals. *Mind*, 488(122):867–913, 2013.
- J.L. Dowell. A flexible contextualist account of epistemic modals. *Philosophers' Imprint*, 11(14):1–25, 2011.
- Dorothy Edgington. Vagueness by degrees. In Keefe and Smith, editors, *Vagueness: A Reader*. MIT Press, Cambridge, MA, 1997.

⁴¹What about non-credal attitude verbs? One option is to generalize Supposes as Support (fn.23). For any non-credal attitude verb N , let N_A^w be the set of worlds compatible with what A N 's at w . Then:

Generalization to Non-Credal Attitude Verbs $s[N_A\phi] = s \cap \{w \mid N_A^w \models \phi\}$.

This predicts that one cannot coherently take any non-credal attitude towards an epistemic contradiction, since no set of worlds supports an epistemic contradiction.

- Andy Egan. Epistemic modals, relativism, and assertion. *Philosophical Studies*, 133(1): 1–22, 2007.
- Andy Egan and Brian Weatherson. Epistemic modals and epistemic modality. In Egan and Weatherson, editors, *Epistemic Modality*. Oxford University Press, Oxford, 2011.
- Andy Egan, John Hawthorne, and Brian Weatherson. Epistemic modals in context. In Preyer and Peter, editors, *Contextualism in Philosophy*. Oxford University Press, Oxford, 2005.
- Jeremy Fantl and Matthew McGrath. *Knowledge in an Uncertain World*. Oxford University Press, Oxford, 2009.
- Richard Foley. *Working Without a Net*. Oxford University Press, 1993.
- Dorit Ganson. Evidentialism and pragmatic constraints on outright belief. *Philosophical Studies*, 139:441–458, 2008.
- Thony Gillies. A new solution to moore’s paradox. *Philosophical Studies*, 105(3), 2001.
- Michael Green and John N. Williams, editors. *New Essays on Moore’s Paradox*. Oxford University Press, New York, 2007.
- Jeroen Groenendijk and Martin Stokhof. Dynamic montague grammar. In Laszlo Kalman and Laszlo Polos, editors, *Papers from the Second Symposium on Logic and Language*, pages 3–48, Budapest, 1990. Akademiai Kiado.
- Jeroen Groenendijk and Martin Stokhof. Dynamic predicate logic. *Linguistics and Philosophy*, 14:39–100, 1991a.
- Jeroen Groenendijk and Martin Stokhof. Two theories of dynamic semantics. *Lecture Notes in Computer Science*, 478:55–64, 1991b.
- Ian Hacking. Possibility. *Philosophical Review*, 76:143–168, 1967.
- John Hawthorne. Eavesdroppers and epistemic modals. *Philosophical Issues*, 17:92–101, 2007.
- John Hawthorne. Knowledge and epistemic necessity. *Philosophical Studies*, 158(3):493–501, 2012.
- John Hawthorne, Daniel Rothschild, and Levi Spectre. Belief is Weak. *Philosophical Studies*, 173:1393–1404, 2016.
- Irene Heim. *The Semantics of Definite and Indefinite Noun Phrases*. PhD thesis, University of Massachusetts, Amherst, MA, 1982.
- Irene Heim. On the projection problem for presuppositions. In *WCCFL*, volume 2, pages 114–125, 1983.
- Irene Heim. Presupposition projection and the semantics of attitude verbs. *Journal of Semantics*, 9(3):183–221, 1992.
- Jaakko Hintikka. *Knowledge and Belief*. Cornell University Press, 1962.
- Jerry R. Hobbs. On the coherence and structure of discourse. Technical Report CSLI-85-37, Stanford University, Center for the Study of Language and Information, 1985.
- Michael Huemer. Moore’s paradox and the norm of belief. In Nuccetelli and Seay, editors, *Themes from G.E. Moore: New Essays in Epistemology and Ethics*, volume 74, pages 142–157. Clarendon Press, Oxford, 2007.
- Lauri Karttunen. Presuppositions and linguistic context. *Theoretical Linguistics*, 1:181–194, 1974.
- Andrew Kehler. *Coherence, Reference, and the Theory of Grammar*. CSLI Publications, 2002.
- Chris Kennedy. The semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy*, 30:1–45, 2007.
- Rodger Kibble. Dynamics of epistemic modality and anaphora. *Proceedings of the International Workshop on Computational Semantics, ITK*, 1995.

- Peter Klecha. Positive and conditional semantics for gradable modals. In Guevara et al, editor, *Proceedings of Sinn und Bedeutung 16*, volume 2, pages 363–376. MIT Working Papers in Linguistics, 2012.
- Nathan Klinedinst and Daniel Rothschild. Quantified epistemic modality, 2015.
- Angelika Kratzer. The notional category of modality. In Eikmeyer and Rieser, editors, *Words, Worlds, and Contexts: New Approaches in Word Semantics*. W. de Gruyter, Berlin, 1981.
- Angelika Kratzer. Modality. In von Stechow and Wunderlich, editors, *Semantics: An International Handbook of Contemporary Research*. W. de Gruyter, Berlin, 1991.
- Angelika Kratzer. *Modals and Conditionals*. Oxford University Press, Oxford, 2012.
- Henry Kyburg. *Probability and the Logic of Rational Belief*. Wesleyan University Press, Middletown, CT, 1961.
- Peter Lasersohn. Pragmatic halos. *Language*, 75:522–551, 1999.
- Hannes Leitgeb. Reducing belief simpliciter to degrees of belief. *Annals of Pure and Applied Logic*, 164:1338–1389, 2013.
- Hannes Leitgeb. The stability theory of belief. *Philosophical Review*, 123(3):131–171, 2014.
- Hannes Leitgeb. *The Stability of Belief*. Oxford University Press, New York, forthcoming.
- David Lewis. Scorekeeping in a language game. *Journal of Philosophical Logic*, 8(1):339–359, 1979.
- David Lewis. Elusive knowledge. *Australasian Journal of Philosophy*, 74(4):549–567, 1996.
- John MacFarlane. Epistemic modals are assessment sensitive. In Egan and Weatherston, editors, *Epistemic Modality*. Oxford University Press, Oxford, 2011.
- David Makinson. The paradox of the preface. *Analysis*, 25(6):205–207, 1965.
- Aidan McGlynn. Believing things unknown. *Noûs*, 47(2):385–407, 2013.
- Aidan McGlynn. *Knowledge First?* Palgrave MacMillan, New York, 2014.
- G.E. Moore. Certainty. In *Philosophical Papers*. Allen and Unwin, London, 1959.
- Sarah Moss. On the semantics and pragmatics of epistemic vocabulary. *Semantics and Pragmatics*, 8(5):1–81, 2015.
- Colin Radford. Knowledge: By examples. *Analysis*, 27:1–11, 1966.
- Craige Roberts. Modal subordination and pronominal anaphora. *Linguistics and Philosophy*, 1989.
- Bertrand Russell. *The Problems of Philosophy*. Henry Holt and Company, New York, 1912.
- Patrick Rysiew. The context-sensitivity of knowledge attributions. *Noûs*, 35(4):477–514, 2001.
- Joshua Schechter. Rational self-doubt and the failure of closure. *Philosophical Studies*, 163(2):428–452, 2013.
- Moritz Schulz. Epistemic modals and informational consequence. *Synthese*, 174:385–395, 2010.
- Roy Sorensen. Meta-agnosticism: Higher order epistemic possibility. *Mind*, 118(471):777–784, 2009.
- Robert Stalnaker. Presuppositions. *Journal of Philosophical Logic*, 2:447–457, 1973.
- Jason Stanley. Fallibilism and concessive knowledge attributions. *Analysis*, 65(2):126–131, 2005.
- Jason Stanley. Knowledge and certainty. *Philosophical Issues*, 18(1):35–57, 2008.
- Tamina Stephenson. Judge dependence, epistemic modals, and predicates of personal taste. *Linguistics and Philosophy*, 30(4):487–525, 2007a.
- Tamina Stephenson. *Towards a Theory of Subjective Meaning*. Ph.D. Thesis, MIT, 2007b.
- Scott Sturgeon. Reason and the grain of belief. *Noûs*, 42(1):359–396, 2008.
- Jonathan Sutton. Stick to what you know. *Noûs*, 39(3):359–396, 2005.
- Jonathan Sutton. *Without Justification*. MIT Press, Cambridge, MA, 2007.

- Peter Unger. *Ignorance*. Clarendon Press, Oxford, 1975.
- Robert van Rooij. A modal analysis of presupposition and modal subordination. *Journal of Semantics*, 2005.
- Frank Veltman. Defaults in update semantics. *Journal of Philosophical Logic*, 25(3):221–261, 1996.
- Kai von Fintel and Thony Gillies. An opinionated guide to epistemic modality. In Gendler and Hawthorne, editors, *Oxford Studies in Epistemology*, volume 2, pages 32–62. Oxford University Press, Oxford, 2007.
- Kai von Fintel and Thony Gillies. *Must... stay... strong!* *Natural Language Semantics*, 18(4):351–383, 2010.
- Brian Weatherson. Can we do without pragmatic encroachment? *Philosophical Perspectives*, 19:417–443, 2005.
- Brian Weatherson. Knowledge, bets, and interests. In Brown and Gerken, editors, *New Essays on Knowledge Ascriptions*, pages 75–103. Oxford University Press, New York, 2012.
- Malte Willer. Dynamics of epistemic modality. *Philosophical Review*, 122(1):44–92, 2013.
- John N Williams. Moorean absurdity and the intentional ‘structure’ of assertion. *Analysis*, 54(3):24–22, 1994.
- Timothy Williamson. *Knowledge and its Limits*. Oxford University Press, Oxford, 2000.
- Alex Worsnip. Possibly false knowledge. *Journal of Philosophy*, 112:225–246, 2015.
- Seth Yalcin. Epistemic modals. *Mind*, 116(464):983–1026, 2007.
- Seth Yalcin. Nonfactualism about epistemic modality. In Egan and Weatherson, editors, *Epistemic Modality*. Oxford University Press, Oxford, 2011.
- Seth Yalcin. A counterexample to modus tollens. *Journal of Philosophical Logic*, 41(6):1001–1024, 2012a.
- Seth Yalcin. Context probabilism. *Proceedings of the 18th Amsterdam Colloquium*, pages 12–21, 2012b.
- Seth Yalcin. Epistemic modality *De Re*. *Ergo*, 2(19):475–527, 2015.
- Seth Yalcin. Belief as question-sensitive. *Philosophy and Phenomenological Research*, forthcoming.
- Henk Zeevat. Presupposition and accommodation in update semantics. *Journal of Semantics*, 9(2):379–412, 1992.