



Still going strong

Kai von Fintel¹ · Anthony S. Gillies²

Accepted: 31 August 2020 / Published online: 6 January 2021
© Springer Nature B.V. 2021

Abstract

In “*Must ...stay ...strong!*” (von Fintel and Gillies in *Nat Lang Semant* 18:351–383, 2010. <https://doi.org/10.1007/s11050-010-9058-2>), we set out to slay a dragon, or rather what we called The Mantra: that epistemic *must* has a modal force weaker than expected from standard modal logic, that it doesn’t entail its prejacent, and that the best explanation for the evidential feel of *must* is a pragmatic explanation. We argued that all three sub-mantras are wrong and offered an explanation according to which *must* is strong, entailing, and the felt indirectness is the product of an evidential presupposition carried by epistemic modals. Mantras being what they are, it is no surprise that each of the sub-mantras have been given new defenses. Here we offer them new problems and update our picture, concluding that *must* is (still) strong.

Keywords Modality · Epistemic · Evidentiality

1 The mantra

In von Fintel and Gillies (2010), we set out to slay a dragon, or rather a mantra about epistemic *must*. The Mantra is a reaction to what we call *Karttunen’s Problem*.¹ There is a two-fold observation at its heart. Here is Karttunen’s example:

¹ The locus classicus is Karttunen (1972). But see also Groenendijk and Stokhof (1975), Lyons (1977), Veltman (1985) and, Kratzer (1991).

We presented this material at the Ohio State Workshop on Modality on March 23, 2016, where our commentator Dan Lassiter provided useful pushback. Comments from students in an advanced semantics class at MIT helped immensely with some of our central arguments. We thank the two reviewers and the editors for incisive and valuable advice.

✉ Kai von Fintel
fintel@mit.edu

Anthony S. Gillies
thonyg@gmail.com

¹ Department of Linguistics and Philosophy, Massachusetts Institute of Technology, Room 32D-808, 77 Massachusetts Avenue, Cambridge, MA 02139, USA

² Department of Philosophy, Rutgers University, 106 Somerset, 5th Floor, New Brunswick, NJ, USA

- (1) a. John left.
b. John must have left.

The first part of the observation is that *must*-claims like (1b) seem to make weaker claims (in some as-yet-unspecified sense of ‘weaker’) than their bare prejacent counterparts like (1a). The literature is full of examples and declarations of intuitions like this.

The second part of the observation is about the surprising contrast in the *evidential distribution of must*: roughly, the better situated a speaker is to the truth of the prejacent the weirder it is for her to issue the corresponding modalized claim and the better off she is using the bare prejacent instead.

- (2) (Billy is looking out the window seeing pouring rain.)
a. It is raining.
b. ??It must be raining.
- (3) (Billy sees people coming in with wet rain gear and knows rain is the only explanation.)
a. It is raining.
b. It must be raining.

The judgment is that in the first case it is fine for Billy to utter (2a) and it is weird for her to utter (2b), while in the second case it is OK for her to utter (3a) and it is OK for her to utter (3b). Again, examples are easy to multiply.

The Mantra is the oft-repeated claim that epistemic *must* is weak. In our telling, this has three sub-mantras:

- W1 *must* is not a strong necessity modal;
W2 *must* doesn’t entail its prejacent; and
W3 the best way to understand the evidential distribution of *must* is by way of a pragmatic derivation of some sort.

It is easy to see how the package of W1–W3 gained its status.² Karttunen’s Problem seems to show that the facts on the ground about *must* are exactly backwards from what you would expect if *must* behaved like a necessity operator quantifying over possibilities compatible with what is known. Meanwhile, W1 and W2 explain the felt weakness of *must*-claims compared to their bare prejacent. With that in hand, the weirdness of things like (2b) compared to (2a) can be explained by straightforward quantity implicature reasoning.

Still, we argued against W1–W3 taken as a package and offered a different solution to Karttunen’s Problem according to which:

- S1 *must* is a strong universal epistemic modal (there are no epistemic modals strictly stronger than it because it is at the top of the scale);

² To be sure, different theories can and do instantiate W1–W3 in different ways.

S2 it entails its prejacent; and

S3 the evidential distribution of *must* (we think this has nothing to do with weakness and has everything to do with indirectness) can be usefully modeled as a presupposition in a semantics that appeals to a more structured characterization of the domain epistemic modals quantify over.

While W1–W3, and so S1–S3, are definitely related, they are in principle separable. New work has taken W1–W3 up in somewhat modular fashion, arguing that: (i) *must* can't be a strong universal epistemic modal (Lassiter 2016; Goodhue 2017); (ii) *must* can't entail its prejacent (Goodhue 2017) or at least our reasons for denying this can't be right (Lassiter 2016); and (iii) pragmatic explanations of the evidential distribution of *must* are possible after all (Goodhue 2017; Mandelkern 2016, 2019). We want to update our opposition accordingly.

Here is what follows. In Sect. 2, we briefly sketch some of the main arguments from von Fintel and Gillies (2010). Then in Sects. 3–5, we address the new generation of responses and defenses. In Sect. 6, we wrap up.

2 Recap

There is some sort of felt contrast between things like (1a) and their modalized counterparts like (1b) but it is way too fast to declare that this felt contrast is a contrast of semantic strength. For all that's been said, it may be that the felt difference is a difference of (in)directness and that may be something compatible with a strong semantics for *must* (indeed, we showed that it is).

So semantic weakness as a diagnosis can't be assumed. But it also leaves a lot unexplained. Given W1–W3 the explanation for why things like (2b) are marked or seriously degraded compared to (2a) can (allegedly) go along straightforward quantity implicature reasoning: if $must\ \phi \not\models \phi$ then (2b) is weaker than expected in the circumstances and so at least misleading. But (3a) and (3b) are both fine. Why would Billy reach for the (allegedly) weaker *must*-claim if she is in a position to assert the (allegedly) stronger bare prejacent? We will return to this in Sect. 5.

Our main empirical point: semantic weakness faces an uphill battle because it doesn't seem to square with a battery of observations about how *must* behaves. Here is a brief rundown.³

Observation 1 Epistemic *must* doesn't always convey weakness.

For instance:

- (4) The ball is in A or in B or in C.
 It's not in A.
 It's not in B.
 So, it must be in C.

³ The examples are all from von Fintel and Gillies (2010).

There's not a whiff of weakness here, but *must* is perfectly at home.⁴

Observation 2 Conjoining *must* ϕ with an expression that ϕ might not be true seems contradictory.

Since W2 explicitly says that *must* $\phi \not\equiv \phi$ we would expect to find it easy and natural and coherent to conjoin *must* ϕ with an expression of the possibility that $\neg \phi$. It doesn't seem to be:

- (5) a. #It must be raining but perhaps it isn't raining.
b. #Perhaps it isn't raining but it must be.

Observation 3 Epistemic *must* doesn't combine well with *only*.

Given W1 and W2 there is space between the top of the scale of epistemic strength and *must*. As with all things not at the top of a scale, we therefore would expect *must* to combine naturally with *only*.

- (6) Alex: It must be raining.
Billy: (opens curtains) No it isn't. You were wrong.
Alex: #I was not! Look, I didn't say it was raining. I only said it **must** be raining. Stop picking on me!

Downplaying what was said with *I only said* is completely impossible with *must*, but substitute weaker expressions (*ought*, *probably*, *might*) and *only* gets along with them like with old friends. This is exactly what one expects from a maximally strong element. We put this point in terms of "distancing" or speaker commitment but that is an optional gloss of the phenomenon.

How, then, to solve Karttunen's Problem without W1–W3? Here is a sketch of the account.

⁴ We noted a similar point for *must* in premises:

- (i) a. If Carl is at the party, then Lenny must be at the party.
b. Carl is at the party.
c. So: Lenny is at the party.

W1 and W2 together predict that (ic) isn't entailed by the premises (ia) and (ib). Our judgment goes the other way. Lassiter (2016, 139–141) is unmoved by this, arguing that this should be explained away: (ic) only seems to be entailed by (ia) and (ib) when it is really only nearly-entailed. Our point is that there is no relevant difference between (i) and (ii):

- (ii) a. If Carl is at the party, then Lenny must be at the party.
b. Carl is at the party.
c. So: Lenny must be at the party.

A theory that says there is a difference but explains it away by insisting that we are systematically mistaken about it is dispreferred to one that embraces the non-difference and predicts it.

Definition 1 (kernels, modal bases). A kernel $K^c(w)$ in a context c at a world w is a non-closed set of propositions, those that are direct (enough) information at w in c . An epistemic modal base $B^c(w)$ in c at w is determined by $K^c(w)$ iff $\bigcap K^c(w) = B^c(w)$.

Note that the relationship between kernels and bases is many-one: distinct kernels can determine the same modal base. For instance (and simplifying a bit), in (2) the kernel is the set of propositions $\{it\ is\ raining\}$ and in (3) it is the set $\{wet\ rain\ gear,\ wet\ rain\ gear\ only\ if\ it\ is\ raining\}$.⁵

Epistemic modals presuppose that the kernel doesn't settle whether their prejacent are true.

Definition 2 (definedness, truth conditions). For any ϕ :

$$\llbracket must\ \phi \rrbracket^{c,w} = \begin{cases} 1 & \text{if } K^c(w) \text{ doesn't settle } \llbracket \phi \rrbracket^c \text{ and } B^c(w) \subseteq \llbracket \phi \rrbracket^c \\ 0 & \text{if } K^c(w) \text{ doesn't settle } \llbracket \phi \rrbracket^c \text{ and } B^c(w) \not\subseteq \llbracket \phi \rrbracket^c \\ \text{undefined} & \text{otherwise} \end{cases}$$

Clearly, the crucial thing is to say when a kernel settles a proposition so that $K^c(w)$ can fail to settle $\llbracket \phi \rrbracket$ even though it entails $\llbracket \phi \rrbracket$. We offered two (not quite equivalent) implementations of this idea.

Definition 3 (settling, two ways). $K^c(w)$ settles P iff ...

- i. ...for some $X \in K^c(w)$: either $X \subseteq P$ or $X \cap P = \emptyset$; or
- ii. ... P is an issue raised by the partition induced by the propositions in $K^c(w)$.

These are not equivalent. But either can be paired with the basic analysis to explain the evidential distribution of *must*, and as a bonus the parallel evidential distribution of *can't*, without resorting to weakness. But, really, any good implementation will work.

3 W1 and S1

W1 and W2 naturally come as a pair: the theories in Kratzer (1991) and Veltman (1985) both treat them that way. And so we did, too. But they can come apart. Here we look at Lassiter's (2016) new responses and defenses of W1 considered on its own merits.

⁵ K for kernel (not knowledge) and B for base (not belief). Formally, kernels are Kratzerian modal bases but we use their structure in a novel way. This definition, like Definition 4 in von Fintel and Gillies (2010), treats all information as either direct enough or as following from what is direct enough. This is a simplification and can be removed. Here's what we said about it before: "It is an optional extra and our story is officially agnostic on it. To remove its trace: introduce an upper bound $U \subseteq W$ representing the not-direct-but-not-inferred information in the context and relativize all our definitions to this upper bound instead" (p. 371).

3.1 Part-time weakness

Consider again Observation 1 and examples like (4).⁶ These show that *must* isn't always weak. But wait: isn't this a "strictly irrelevant" point? After all, "no 'Mantrista' has claimed that *must* entails a lack of certainty" (Lassiter 2016, 137). There are several ways to think about strong uses of weak items. Let us discuss three such ways.

First way: *must*-claims are always semantically weak but (of course) they are still true in "strong" scenarios like (4). Kratzer (1991) glosses the semantic weakness of *must* this way: "In uttering [something like] (3b) rather than (3a), I convey that I don't rely on known facts alone." Our point is that this is a weird signal to send your hearer when you deduce the ball's location. A parallel: *every pizza joint in town* quantifies over all the pizza joints in town while *the best pizza joints in town* quantifies over a privileged subset of them delivered by some ranking.

- (7) a. The best pizza joints in town serve wine.
b. Every pizza joint in town serves wine.

Suppose you know (7b) is true. Then it is weird and at least misleading to utter (7a) even though it, too, must be true.

Second way: *must* is sometimes semantically weak and sometimes semantically strong. This can happen, for instance, with Kratzer's weak *must*: it is a universal quantifier over the most normal worlds compatible with an epistemic modal base where normalness is induced by the propositions in an ordering source. So, if the ordering source is empty (or if all the possibilities in the modal base are tied or incomparable according to it) then *must* ends up strong after all, quantifying over all the worlds in the modal base. This is not quite compatible with what Kratzer says about there being a signal when using *must*, but set that aside.

We (still) acknowledge the formal point but deny that this shows that these uses of *must* are "strictly irrelevant". For starters, hearers would have to be sure that the weak *must* is being used since otherwise the signal of weakness would be lost on them. Beyond that, it is also a bit weird to say that *must* can use an ordering but doesn't have to. This would make it unlike its quantificational cousin *the best*. Suppose we stipulate that no meaningful comparisons can be made among the pizza joints: even so, we can't use (7a) to mean (7b). Why would *must* be so different from *the best*? Probably it isn't.

Third way, in fact the one argued for in Lassiter (2016): *must* is weak but has no relevant strictly stronger competitor that would be preferred in a strong use situation. Lassiter agrees with us that *must* has an evidential component and any conceivable stronger competitors such as *know* or *certain* do not. So, a speaker could choose weak *must* in a strong use case to highlight the subtle nuance of evidentiality. We do not know what precisely the underlying theory of pragmatic competition would be like and so it could be that this approach will be able to deal with our cases. But dealt

⁶ It has been suggested that maybe this *must* isn't an epistemic *must* but alethic, truth-in-all-possible-worlds *must* (Giannakidou 1999; Goodhue 2017; Giannakidou and Mari 2018). The question is: why not epistemic? If the answer is because there is no weakness, then the suggestion is unmotivated. If the answer is instead more along the lines of the argument we consider in the main text, then that would be principled but, as we'll argue, still not convincing.

with they must be and hence they aren't strictly irrelevant. Let us also note that when inferential evidentiality is marked explicitly with *therefore* or *so*, *must* is still just as good as the plain prejacent or *know/certain*, and so its choice is presumably not just motivated by the signal of evidentiality.

We conclude that uses of *must* like those in (4) aren't "strictly irrelevant". They point to the fact that you can attach *must* with reckless abandon to the conclusion of any valid argument, without any hint of implicated weakness. This is unexpected given W1.

3.2 Interaction with *only*

We claim that *must* and *only* don't mix well (Observation 3) and that you'd expect otherwise if W1 were on the right track. Against this, Lassiter reports that *only* doesn't go with expressions that are merely *near-maximal* either. This could be undermining, if true. His example:

- (8) Alex: It's 99.9% certain that it's raining.
 Billy: (*opens curtains*) No it isn't. You were wrong.
 Alex: ??I was not! Look, I didn't say it was raining. I only said it was 99.9% certain that it was. Stop picking on me!

We do not share the judgment that Alex's rejoinder is marked or degraded. In any case, there are replies open to Alex that mix *only* and *99.9% certain*:

- (9) Alex: It's 99.9% certain that it's raining.
 Billy: (*opens curtains*) No it isn't. You were wrong.
 Alex: Well, strictly speaking, I was not wrong. I was careful. I only said it was 99.9% certain that it was raining.

This now seems definitely fine. And this kind of adjustment doesn't rescue the combination of *only* and *must*:⁷

- (10) Alex: It must be raining.
 Billy: (*opens curtains*) No it isn't. You were wrong.
 Alex: Well, strictly speaking, I was not wrong. I was careful. #I only said it must be raining.

The essential observation is that in the context of a scale *only* doesn't go with items at the top of the scale and goes with non-maximal elements instead. Consider a quantificational parallel:

⁷ One reviewer disagreed with our judgment and found no significant contrast between (9) and (10). It is tempting to despair over such disagreements between speakers with vested interests. We note that Del Pinal and Waldon's (2019) experiments corroborated that there is a significant contrast in this case.

- (11) Alex: All/most/many/some student(s) are from abroad.
 Billy: Hey, Naomi isn't. So, you're wrong.
 Alex: I was not! Look, *I only said (#all/most/many/some) students are from abroad.*

This is precisely what we see with modals.⁸

4 W2 and S2

Our solution to Karttunen's Problem goes for S2 and against W2, embracing the thesis that *must* ϕ asymmetrically entails ϕ . The mechanism we defended treats *must* as a universal quantifier over $B^c(w)$. Since you can't know what isn't true, it then follows for every w that $w \in B^c(w)$ and so that *must* $\phi \models \phi$. The thesis and the mechanism are related but, again, separable. And each has been on the receiving end of criticism in Lassiter (2016) and Goodhue (2017).

4.1 Problematic "conjunctions"

If *must* ϕ didn't entail ϕ then *must* ϕ would be compatible with expressing that ϕ just might not be true. But the flat-footed conjunctions in (5) are unacceptable. That is the general pattern reported in Observation 2.

However (we conceded) there is a loophole: perhaps *perhaps* isn't a weak existential modal but is the dual to (allegedly) weak *must*. In that case the contradictoriness of (5) is predicted by all parties since *must* $\phi \models \neg \textit{perhaps} \neg \phi$. We said: fine, just find a weak existential modal expressing epistemic possibility and it will be horrible in place of *perhaps* in (5). (We will return to the loophole below.)

Maybe we were wrong. Looking at corpus data, Lassiter reports naturally-occurring passages that feature a juxtaposition of a *must*-claim and an explicit expression of speaker uncertainty.⁹ We admit that the examples strike us as more or less felicitous — at least, we'll treat them that way to give ourselves a bigger burden.¹⁰

⁸ Lassiter (2016, 143) argues that the unacceptability of (i) undercuts Observation 3:

- (i) Alex: ??I was not! Look, I didn't say it was raining. I only said I was absolutely certain that it was. Stop picking on me!

This corroborates what we've been saying: *absolutely certain* is a top element in the scale of certainty and so we would not expect it to combine with *only*.

⁹ That's not quite what we promised wouldn't happen. In any case, we (of course) welcome data wherever it can be found: the lab, the wild, and the armchair. What we don't agree with is some of the surrounding rhetoric that the examples collected from the wild are somehow more probative than intuitive judgments about homegrown examples because the latter are "mere intuitions" (Lassiter 2016, 139). In the end Lassiter relies on his judgment that the examples from his corpus search are coherent and sensible and invites us (collectively) to share that judgment. So it's "mere intuitions" all around or nowhere. We prefer to call it all data. And as we discuss in the text, we think there is a clear and motivated explanation for such data compatible with S1–S3.

¹⁰ Del Pinal and Waldon (2019) provide further corroboration that these conjunctive passages are acceptable to some significant degree.

There are two kinds of examples to think about. First, there are examples where a *must*-claim is followed by an expression of uncertainty or lack of sure knowledge:

- (12) This is a very early, very correct Mustang that has been in a private collection for a long time. ... The speedo[meter] shows 38,000 miles and *it must be* 138,000, but *I don't know for sure*.
- (13) I have an injected TB42 turbo and don't like the current setup. There is an extra injected located in the piping from the throttle body.. *Must be* an old DTS diesel setup but *I'm not certain*. Why would they have added this extra injector?

Secondly, there are examples where rumination that shows uncertainty is nevertheless concluded with a *must*-claim:

- (14) I refuse to believe that this one game, Lost Planet 2 DX11, which was previously 100% stable remember, is crashing because my overclock is unstable It's *not impossible*, granted, but IMO it is *highly unlikely*. *There must be* some other cause.

We should have known better than to make the rash promise we made, since we know all too well about the rapidly shifting grounds of modal conversation (von Fintel 2001; Gillies 2007). We think that is what is going on here, too: Lassiter has found some examples in the wild where epistemic modals undergo shifts in the possibilities deemed relevant, the modal horizon. In the first sort of example, a *must* is followed by uncertainty: this is a speaker who is expanding the modal horizon. In the second sort of example, a speaker is considering some possibilities but then concludes with a *must*: here there's a decision to reset the modal horizon to a more realistic boundary. Both sorts of cases involve instability across contexts and the claim that *must* ϕ entails ϕ is about what happens within a given context.

Does the kind of wriggle room we're allowing ourselves completely insulate our analysis from counterexamples? In other words, are we putting the S1–S3 package on “pragmatic life support“?¹¹ Or, is there a way to reliably prevent shifts so that one can observe the predicted unacceptable conjunctions?

We think that there are two ways to see that we're correct about these cases. They both offer ways of controlling for contextual instability: if (and only if) the acceptability of (12)/(13) and (14) is due to shifts in modal horizon then speakers shouldn't be willing to embrace both conjuncts simultaneously.¹²

First, there are natural questions we can ask a speaker who goes from *must* to acknowledging uncertainty/lack of sure knowledge:

¹¹ As Dan Lassiter put it at a workshop at the Ohio State University.

¹² To forestall confusion: this is not the same thing as saying that speakers shouldn't be willing to utter or assent to such conjunctions. We'll return to this point below.

- (15) A: That must be an old DTS diesel setup but I'm not certain. Why would they have added this extra injector?
 B: So, given that you're not certain, do you still think that it must be an old DTS diesel setup?
 A: I guess not./Yeah, it must be; I'm sure of it./??Like I said: it must be but I'm not certain.

There are choices: A will either have to walk back the *must*-claim or reset the modal horizon by excluding the additional possibilities. But sticking to both conjuncts doesn't seem open.¹³

Secondly, we can explore the acceptability of these problematic conjunctions in environments that naturally limit shifts in context and thus shifts in the modal horizon. So if the problematic conjunctions are problematic in these environments, this is evidence that their acceptability depends on the availability of context shifts. Here we mention two such environments: *although*-prefixes and embeddings.¹⁴

These examples involve a speaker reporting some reasoning, considering and rejecting possibilities. If we set up a statement that comes after such a reasoning process and states its conclusion, we find it much harder to allow the kind of conjunction that's at stake.¹⁵

- (16) a. #Although I'm not certain, it must be an old DTS diesel setup.
 b. #Although I don't know for sure, it must be 138,000.

As for embedding, consider this scenario: to establish whether a patient has the disease, there are two tests. Test A is cheap, but not always definitive: it can indicate that the patient has the disease but often it merely indicates that a patient is more or less likely to have the disease. There's a second, always definitive test (Test B), which is hugely more expensive. So, the insurance company has rules like (17a) but not like (17b):

- (17) a. Test B can only be administered if the results of Test A are that it is not certain that the patient has the disease but that she likely has it.
 b. #Test B can only be administered if the results of Test A are that it is not certain that the patient has the disease but that she must have it.

And in the insurance company's training manual, you might find a question like (18a) but not (18b):

¹³ We note that both reviewers disagreed with our judgment that there is a difference in acceptability in A's replies.

¹⁴ The *although* device is borrowed from an interesting argument by Kroch (1974, 190–191), who used it to show that definite plurals, even though in some sense they allow exceptions, behave like universal quantifiers in controlled conjunctions:

- (i) a. #Although the townspeople are asleep, some of them are awake.
 b. Although more or less all the townspeople are asleep, some of them are awake.

See Laserson (1999, 523) and Križ (2015, 4) for some discussion.

¹⁵ Again, there is disagreement about this judgment from the reviewers.

- (18) a. Suppose the results of Test A are that it is not certain that the patient has the disease but that she likely has it. Should we approve Test B?
- b. #Suppose the results of Test A are that it is not certain that the patient has the disease but that she must have it. Should we approve Test B?

So: it looks like the acceptability (such as it is) of a juxtaposition of *must* ϕ with an expression that assigns non-zero possibility to $\neg\phi$ is due to contextual shifts in the modal horizon.¹⁶ On the flip-side: when such shifts are blocked, the conjunctions are unacceptable and contradictory-feeling.¹⁷

Lassiter (2016, 139) remarks that in a many-subject experiment describing a lottery drawing (more on which below) 58% of the subjects agreed with a *must* $\neg\phi$ prompt and 92% agreed with a *ϕ is possible* prompt.¹⁸ From this he concludes that “half or more ... were simultaneously inclined to endorse” both (or, as we say, their conjunction).

The conclusion is too fast. The experiment used a between-subject design and so can't really probe the extent to which people want to simultaneously endorse *must* $\neg\phi$ and *ϕ is possible*. Between-subject designs are fine, of course: often, we can infer “speakers (full stop) assent to both *X* and *Y* simultaneously” from “speakers in *A* assent to *X*” and “speakers in *B* assent to *Y*”. But two things are required to allow the inference to bridge the gap in the experiment reported: (i) the events of assenting to *X* and assenting to *Y* must be independent, and (ii) it has to be that assenting to *X* and assenting to *Y* is a fool-proof mechanism revealing that speakers judge both items to be jointly possible. In this situation, we think neither condition is met: part of what is at stake in our dispute is whether the conjuncts in problematic conjunctions are independent and, as we've just seen, assenting to a conjunction need not be a reliable guide to whether, in the final analysis, speakers want to simultaneously get behind both conjuncts. Since we don't have both (i) and (ii), the reported many-subject experiment doesn't speak to whether *must* $\neg\phi$ and an expression that $\neg\phi$ might not be true can coherently hang together.¹⁹

Finally, let's revisit the original loophole: that *perhaps* and *maybe* and *might* are duals to (allegedly) weak *must* and thus not weak but strong existential modals. This would adequately explain the unacceptability of (5). So assume, for the sake of argument, that *it's possible* is weaker than *perhaps/maybe/might*. This leads to a dilemma that turns on two observations about epistemic *can't*.

The first observation we pointed out in von Fintel and Gillies (2010, 373):

Observation 4 Epistemic *can't* patterns like *must* in its evidential distribution.²⁰

16 There are coherent readings of *x is not certain that ϕ but must ϕ* where there are multiple bodies of information (one for *x* and one for us). But because of the multiplicity of bodies of information such readings don't speak directly to whether *must ϕ* is compatible with expressions that ϕ might not be true. Note that this kind of differential targeting isn't possible in (17) and (18).

17 For once, it is pleasing to note that both reviewers agree with our judgments in these cases.

18 We borrow the ‘many-subject’ vs. ‘few-subject’ terminology from Jacobson (2018).

19 Our assessment has since been confirmed in an experiment reported at the 2020 CUNY Sentence Processing conference by Ricciardi et al. (2020), who found that in a within-subject design, Lassiter's results do not persist.

20 A reviewer finds that *can't* is stronger than *must not*. We don't share that judgment.

Some examples:

(19) (Billy looking out the window seeing brilliant sunshine)

- a. It isn't raining.
- b. ??It can't be raining.

(20) (Billy seeing people coming in with sunglasses and parasols and knowing sunshine is the only cause)

- a. It isn't raining.
- b. It can't be raining.

As with *must*: the modalized report is weird when she is looking out the window but modalized or non-modalized will do when she is seeing people putting away their sungear.

The second observation:

Observation 5 Epistemic *can't* ϕ seems incompatible with *it's possible that* ϕ .

Billy's reply in (21b) is fine and to the point:

- (21) a. Alex: Hey, is it possible the keys are in the drawer?
 b. Billy: No, they can't be.

Billy is not leaving the door open just a little about whether the keys are in the drawer: she is explicitly denying that and closing off that possibility. Similarly, the suppositions in (22) are incoherent.²¹

- (22) a. #Suppose it's possible the keys are in the drawer but they can't be.
 b. #If it's possible the keys are in the drawer but they can't be, then ...

Now the dilemma: is *can't* the negation of a strong existential modal (e.g., the allegedly strong *perhaps*) or a weak existential modal (e.g., *it's possible*)? Assuming W2, neither horn covers both Observation 4 and Observation 5. If *can't* is the negation of a strong existential, then it expresses weak necessity. Thus we could (maybe) explain the evidential distribution of *can't*. But then we can't explain the response in (21): Billy's reply would miss the mark and, indeed, she could reply this way:

- (23) a. Alex: Hey, is it possible the keys are in the drawer?
 b. Billy: #Yes, but they can't be.

Which, as a reply, is ...whoa. Similarly, if *can't* is weak then the suppositions in (22) would be mundane and coherent. They aren't and so it can't be.

If, on the other hand, *can't* is the negation of a weak existential then it expresses strong necessity and the explanatory gaps reverse: we would have an explanation of (21) and (22) but not of (20) and (19).

This is all bad news for W2 and indicates, as S2 claims, that *can't* $\phi \models \neg\phi$ and, indeed, that *must* $\phi \models \phi$.

²¹ For the record: we have the same judgments in (21) and (22) with *it's possible* replaced by *perhaps/maybelmight*.

4.2 Anti-knowledge

We now turn to the first of two criticisms of the mechanism that a *must* (in c , at w) involves quantifying over $B^c(w)$. The argument comes from Goodhue's defense of the *epistemic account*, according to which *must* ϕ is felicitous only if ϕ is not known.²²

Here are the trouble-making examples (with extensive context reproduced verbatim).

- (24) (Phil is cooking chicken and peas for his family. When the timer goes off, he checks the chicken's temperature and discovers it is done. He tastes the peas and they are also ready. The table is already set.)
- a. Phil's daughter: Is dinner ready?
 - b. Phil: # Dinner must be ready.
- (25) (Phil is cooking dinner for his family and his friend Meryl. He had to step out in a hurry, and instructed Meryl as he left: "Please turn the peas off when they are done, and take the chicken out of the oven when the temperature is right." When the peas are done, Meryl turns the burner off, and when the chicken is done, she removes it from the oven. She has also seen that the table is set. She wonders whether Phil was planning to make anything else, for example a salad, but Phil didn't mention anything.)
- a. Phil's daughter: Is dinner ready?
 - b. Meryl: Dinner must be ready.

Whether this is problematic depends on whether the kernel $K^c(w)$ settles whether dinner is ready.²³ The allegation is that the theory faces a dilemma: either provide an analysis of what counts as "direct (enough) information" in a context or don't. On the one hand, plausible analyses of "direct (enough) information" are likely to deliver the same verdicts about (24) and (25), and hence will not predict the contrast between

²² There is a similar worry expressed in Sherman (2018).

²³ We note that a '#' in (24a) is a little harsh, since minor adjustments can help a lot. For instance:

- (24') (Same story as before)
- a. Phil's daughter: Is dinner ready?
 - b. Phil: The chicken is done, the peas are ready, and the table is set. So: Dinner must be ready.

Similarly, we're not so sure Meryl is in the clear in (25b). For instance, if she vocalizes her uncertainty about a salad:

- (25') (Same story as before)
- a. Phil's daughter: Is dinner ready?
 - b. Meryl: I don't know whether we are supposed to have salad, too. ??Dinner must be ready.

(24) and (25).²⁴ On the other hand, leaving “direct (enough) information” unanalyzed makes the theory too imprecise to be useful, and hence on that score can’t predict the data. Worse, there is a nearby and clear competitor: the difference is that Phil knows that dinner is ready and Meryl doesn’t.

We will take up the idea of “direct (enough) information” first and then turn to the competitor of invoking anti-knowledge.

The important thing for the theory we defended is that the direct (enough) information in a context doesn’t settle the preajcent. An example: imagine that you’re making a complex dish from a Rick Bayless recipe.²⁵ When you’re a novice, you follow the recipe step by step, checking things off as you accomplish them. Later, when you have full command of the dish, you follow the recipe in a more vague way relying instead on your sense of things. Now imagine that on two occasions (one when you are a novice and one when you are in full command) you see and respond to the stages in the cooking — your “perceptions” and your actions prompted by them — in exactly the same way. In the novice episode you see that the sauce has been simmering for the time prescribed and that it now has the consistency described by the recipe. In the expert episode, you just see that it is done.

- (26) a. (Novice episode)
It must be ready./It’s ready.
b. (Expert episode)
??It must be ready./It’s ready.

The difference is in what’s direct (enough). So the extent to which you think Phil’s reply in (24b) is marked is exactly the extent to which you think that he has direct enough information about what is for dinner and thus direct enough information about whether all the to-dos have been accomplished. Similarly, the extent to which you find Meryl’s reply in (25b) felicitous and true tracks the extent to which you think her information about what is for dinner is not direct enough.

What counts as direct (enough) information is context dependent and imprecise. This doesn’t thereby make a theory that invokes it hopeless and not explanatory. The reverse is true: acceptable uses of *must* are context dependent and imprecise in exactly the same ways that “direct enough information” is. The two things flex and bend together.²⁶

Consider, sorites-style, a heap of sand. Every minute on the minute we remove a single grain. Now, imagine being presented with a collection C_i of the grains out of context so you don’t know where in the grain-removal process C_i lies. No matter the size of the collection, (27) is weird.

²⁴ Goodhue talks about a theory according to which having “identical perceptions” means having the same direct enough information. We take it this is just an example of a (doomed) analysis of direct enough information.

²⁵ <http://www.rickbayless.com/recipe/classic-red-mole/>.

²⁶ Compare von Fintel and Gillies (2010, 370, fn.29).

(27) It must be a heap.

An explanation: if C_i is toward either end of the spectrum then either it is certainly and obviously a heap or certainly and obviously not a heap and the direct enough information settles this. Same thing goes if it is in the interesting middle where the whole point is that the direct enough information doesn't settle things. Adding some context makes a difference, but there are still borderline cases. If you are presented with any two adjacent collections C_i and C_{i+1} , then using (27) to describe C_{i+1} is only as OK as using the bare prejacent *It_i is a heap* is OK to describe C_i . Thus there is an explanation for both the clear-cut cases and the fuzzy middle ground, too.

What would a from-first-principles analysis of “direct enough information” in a context look like? Hard to say, but it would be ill-suited to do the job that is carved out for it. So it's no mark against a theory that it doesn't start by giving a characterization, once and for all, of what counts as direct enough information.

There is a suggestive parallel in the case of counterfactuals and orderings of similarity that get implicated in their semantics. Here is Lewis (1979), arguing that his “Analysis 2” should not be tested against putative pre-theoretical intuitions about what counts as similarity:

The thing to do is not to start by deciding, once and for all, what we think about similarity of worlds, so that we can afterwards use these decisions to test Analysis 2. What that would test would be the combination of Analysis 2 with a foolish denial of the shiftiness of similarity. Rather, we must use what we know about the truth and falsity of counterfactuals to see if we can find some sort of similarity relation —not necessarily the first one that springs to mind —that combines with Analysis 2 to yield the proper truth conditions. It is this combination that can be tested against our knowledge of counterfactuals, not Analysis 2 by itself. In looking for a combination that will stand up to the test, we must use what we know about counterfactuals to find out about the appropriate similarity relation —not the other way around. (Lewis 1979, 466–467)

As with counterfactuals and similarity, so too with epistemic modals and direct enough information: sometimes it is what we know about the truth, falsity, and appropriateness of *must*-claims in a context that will shed light on what information is direct enough.

The competitor thesis that *must* ϕ is felicitous only if ϕ is not known is bold, wildy so. Some examples:

(28) (Billy seeing people with wet rain gear and knowing rain is the only cause)

Billy: It must be raining.

Alex: ??So you don't know it is raining.²⁷

Billy: ??Right, like I said, it must be raining.

This is not a normal conversation. On the flip-side of things:

²⁷ A reviewer suggests that with the addition of *for sure*, Alex's response is improved. This may be so, but Billy's sticking to her guns is still off.

- (29) Alex: Do you know if it's raining?
 Billy: Yes, it must be [because of those wet umbrellas]

This is a completely normal conversation: Billy's *yes* and *must* are both appropriate. That is squarely at odds with the thesis that *must* ϕ is felicitous only if ϕ is not known. That thesis predicts that Billy's reply is some sort of pragmatic contradiction: since her *must* ϕ is OK, it follows by the thesis that she doesn't know ϕ and so her *yes* is out.²⁸

4.3 Knowledge

That *must* requires anti-knowledge is a bit much, but maybe it is also a bit much to say that *must* requires knowledge. And, indeed, results of a many-subject experiment put pressure on this (Lassiter 2016). Participants read the following story:

Yesterday, Bill bought a single ticket in a raffle with 1000 total tickets. There were also 999 other people who bought one ticket each. That is, the tickets were distributed like this:

People holding one ticket: Bill, Mary, Jane, ... [997 more]

The drawing was held last night, and the winner will be announced this evening.

They were then shown a sentence and had to choose between "Agree" and "Disagree".

We highlight the following results: (i) a bare assertion *Bill did not win* is accepted by (slightly) more participants than the *must*-claim *Bill must not have won*; (ii) expressions of knowledge (*We know that Bill did not win*) and certainty (*It is certain that Bill did not win*) are accepted less frequently than the *must*-claim. Note that the data are compatible with our central claim that *must* ϕ entails ϕ .

Lassiter argues that a strong epistemic necessity account of *must* is committed to an entailment from the *must*-claim to knowledge and certainty claims and that therefore these (many-subject) experimental results are a refutation of that semantics for *must*. We disagree.²⁹ What we see is evidence of a difference in sensitivity: the subjects are sensitive to a 0.1% chance of Bill's winning when it comes to judging the knowledge ascriptions and certainty ascriptions (which are quite reliably rejected) but they are not

28 Above we noted that the spectrum of non-borderline and borderline uses of *must* makes sense. Those reasons also cut against the thesis that *must* is only felicitous if the prejacent is unknown. Here's why. Presented with collection C_{i+1} out of the blue, it's definitely weird to say *It must be a heap*. But being presented with C_{i+1} and being told that C_i is a heap, it is fine to say that C_{i+1} must be a heap. If *must* required anti-knowledge it would then follow that while you know that C_i is a heap (because you were told), you don't know that C_{i+1} is. This is exactly backwards from what makes vagueness hard: all those little bridge conditionals *If C_i is a heap then C_{i+1} is a heap* between adjacent collections seem obviously true (and known).

29 We could at this point insist that our version of a strong *must* does not in fact say that *must*-claims entail knowledge and certainty claims. Our gloss of *must* carefully used impersonal phrases such as *it follows from the information that* or *worlds compatible with what is known*. This is because of the widely known (but largely orthogonal to Karttunen's Problem) feature of epistemic modality that it isn't constrained to be speaker-ego-centric. We have commented on this phenomenon in our other work on epistemic modals (in particular von Fintel and Gillies 2007 and von Fintel and Gillies 2011). This interacts in interesting ways with the fact that *must* entails its prejacent, but we set these issues aside here.

sensitive to that same chance when it comes to judging the *must*-claim. It's tempting to think that this shows that knowledge and certainty claims are semantically stronger than *must*-claims. Lassiter gives in to this temptation, but we think it can be resisted for principled reasons.

The general pattern is that speakers make some strong claims and then back off them a bit, or speakers hesitate in the presence of a salient chance of error from making those claims. An *NALS* reviewer commented (on our original paper): "... one might then naturally conjecture that this weaker epistemic position is reflected in the truth conditions of sentences used to communicate information possessed only indirectly" (von Fintel and Gillies 2010, p. 362, fn.19). If there were such a process, we would find a lot of instability: a language tries to have expressions with a strong semantics but speakers who use that expression are inferred to be on shakier ground than those who use more cautious expressions and so the semantics of the strong expression becomes weaker by reflecting the weaker epistemic position. If so, why hasn't "every" evolved to mean "almost every"?

One of the most important results of natural language semantics is that the meanings that travel in conversations are in a very complex way constituted from semantic encoding in multiple dimensions: truth-conditional at-issue content, presuppositions, extra-dimensional conventional implicatures, conversational implicatures, and other highly situational inferences. Famously, Grice (1967) argued that one should be careful about conflating dimensions of meaning and attributing to truth-conditional at-issue content what can and should be seen as coming from other aspects of meaning. We declare that in addition to Grice's "Modified Occam Razor", there needs to be a corollary razor:³⁰

Constraint (New Razor). Do not weaken semantics beyond necessity (i.e. just because there's some speaker uncertainty)!

The quick reaction of reaching for a weak semantics to account for the interesting and varied uses of *must* is the mistake of a theory being too one-dimensional.

That's the why, here's the how. Modal claims and knowledge/certainty ascriptions are not the same thing. So even assuming that *must* ϕ entails *it is known/certain that* ϕ it doesn't follow that *must*-claims and knowledge/certainty ascriptions can be made in all and only the same situations. They may have different pragmatic sensitivities. We suspect this is the case and conjecture that epistemic modals and knowledge ascriptions have different slack tolerance. This would not be unprecedented: we know there are truth-conditionally equivalent (or relevantly similar) expressions that differ in their allowance for slack.

For instance: the difference between definite plurals and explicitly universal quantification.

- (30) a. The villagers are asleep.
 b. All the villagers are asleep.

Notoriously, (30a) allows some of the villagers to be awake, as long as that fact doesn't affect the rhetorical point of the utterance. The explicit universal quantification

³⁰ We like to call this "Shatner's Razor" for reasons we're happy to reveal over a drink.

in (30b), on the other hand, does not tolerate any such exceptions. But this isn't a difference in truth conditions: if Alex utters (30a) and Billy interjects *Wait, the baker isn't!*, it's natural for Alex to concede that strictly speaking not all of the villagers are asleep.³¹

The idea is that epistemic modals and knowledge-ascribing language are similarly related. Both quantify over the same sorts of possibilities, but knowledge ascriptions are less slack tolerant: it's harder to ignore not-ruled-out possibilities for them and a bit easier to ignore them for modal claims. And of course once you make those possibilities explicit, neither construction can ignore them any longer. We think that epistemic modal claims are in a way less explicit and more underspecified than explicit knowledge claims. This allows them to traffic in a lot of *misdirection* and that misdirection gets profitably exploited in conversation (von Fintel and Gillies 2008, 2011). Our willingness to sometimes ignore some $\neg\phi$ possibility when it doesn't get too much in the way of a *must*-claim may be more of the same.

5 W3 and S3

Finally, we consider W3 and S3: what is the best explanation for the evidential distribution of *must*? In von Fintel and Gillies (2010) we encoded *must*'s evidential signal as a presupposition: $\llbracket \text{must } \phi \rrbracket^{c,w}$ is defined only if $K^c(w)$ doesn't settle $\llbracket \phi \rrbracket^c$. We admitted that the signal of indirectness seems crosslinguistically robust and so in some ways an implicature-based explanation would be desirable. Readers clearly heard us. But what wasn't so clear to them is that there seems at present no implicature story that is both sufficiently detailed to provide an explanation and that covers what needs covering. Meanwhile, a presupposition-based explanation can be given that is both explanatory and empirically supported.³²

5.1 Quantity

Suppose asserting ϕ somehow conveys to your audience that you know ϕ .³³ Then you might suspect that more or less standard quantity implicature reasoning can predict the evidential distribution of *must*. A new version of this sort of explanation is defended in Goodhue (2017).

Here again is the canonical pattern that needs explaining:

- (2) (Billy is looking out the window seeing pouring rain.)
- a. It is raining.
 - b. ??It must be raining.

31 We are therefore attracted to analyses such as the one in Lasersohn (1999) (but also, for example, Križ 2015) that locate the difference outside the truth conditions.

32 A reviewer notes that Ippolito (2018, 610–611) attempts to show that at least in the case of *might*, the evidential presupposition does not project as our theory would predict. We acknowledge that there is more to be said.

33 Whether this is achieved by invoking knowledge as the norm of assertion or some other way doesn't matter for us.

- (3) (Billy sees people coming in with wet rain gear and knows rain is the only explanation.)
- a. It is raining.
 - b. It must be raining.

Assuming W1 and W2, there seems to be a pragmatic explanation. Since Billy (in c , at w) is in a position to assert (2a), her information (the worlds in $B^c(w)$) includes only raining worlds. This asymmetrically implies the truth conditions for *It must be raining*. Hence if her information includes only raining-worlds, it is misleading to assert (2b). The wrinkle that Goodhue adds is that the evidential signal of *must* isn't about directness, but about whether the prejacent is known.

The problem with this explanation is that when Billy's information is indirect, her choice is unconstrained: both (3a) and (3b) are fine. Since Billy *can* use the bare prejacent here, then the pragmatic derivation on offer predicts that she *has to* use the bare prejacent. This is at odds with the fact that the modal is fine, too. Even worse: this unconstrained feature is ubiquitous. We conjecture that if a *must* is OK (and we control for contextual instability), an assertion of the bare prejacent is also OK. So, assuming a quantity implicature implementation of W3 along these lines, we get the uncomfortable prediction that anytime you can use a *must* you can't. This is a fully general problem for any explanation of the evidential distribution of *must* that pairs semantic weakness with quantity implicature reasoning. Such a defense must either argue that speakers aren't unconstrained when their information is indirect or else explain why we have *must* at all.

5.2 Non-quantity

If quantity isn't the right place to look, maybe we can look elsewhere. This would mean pairing W3 with S1 and S2. This has been pursued in Mandelkern (2019) (based on Mandelkern 2016).

The broad outline of the explanation runs like this. First, the signal of indirectness is ultimately derived from a novel pragmatic constraint: that an utterance of *must ϕ* is felicitous only if there is a mutually salient argument for ϕ . The idea is that it is this constraint that can serve as a basis for a pragmatic derivation of the signal of indirectness. For that, a few more pieces are required. So, second, ϕ and *must ϕ* are genuine alternatives to each other and the characteristic effect of an utterance of *must ϕ* is to (propose to) update the common ground with $\llbracket\phi\rrbracket$ on the basis of a shared and mutually available argument.³⁴ Third, it is bad to draw people's attention to mutually available arguments that either too obviously support ϕ or are not your best information that ϕ , and that is why an utterance of *must ϕ* signals that the speaker's

³⁴ Mandelkern goes on to argue that the novel pragmatic constraint requiring a mutually salient argument is itself amenable to a pragmatic derivation (via a manner implicature) and that the argument for that predicts that *must ϕ* amounts to a proposal to add $\llbracket\phi\rrbracket$ to the common ground on the basis of a shared and mutually available argument. As an aside, we are skeptical about tying the upshot of *must ϕ* so closely to trying to coordinate everyone on ϕ . Some uses of *must ϕ* manifestly do not have this coordinating effect but even in such uses the evidential signal of indirectness remains.

best information about ϕ is that it follows not too obviously (i.e., indirectly) from this mutually available argument.

The major load-bearing in all of this is done by the novel pragmatic constraint that *must* requires a mutually salient argument for its prejacent. We will focus on that.

Mandelkern provides prima facie evidence for such a constraint:

- (31) (Patch the rabbit sometimes gets into the box where her hay is stored. On his way out, Mark hears a snuffling from the box. At work, Bernhard asks him how Patch is.)
- a. She's great. She got into the hay box this morning.
 - b. She's great. She must have gotten into the hay box this morning.

If that's all there is to the conversation, (31b) is a little weird. But the *must* is fine if Mark had first (or perhaps right after) said *I heard a snuffling from the box of hay on my way out*. The judgment is subtle, but we agree there is a preference for the non-modal in (31). This is, of course, compatible with the view defended in von Fintel and Gillies (2010): there we insisted that an utterance of *must* ϕ is felicitous only if ϕ isn't settled by the kernel encoding the direct enough information. This is officially silent on the question as to whether there are additional constraints that must be met in certain situations.

For now we only want to take a stand on whether the proposed additional constraint about a mutually salient argument is enough to ground indirectness and hence be a crucial cog in explaining the evidential distribution of *must*. The short answer: no.

Unlike the signal of indirectness, the proposed signal about a mutually salient argument isn't always present. Thus we can't expect to derive indirectness from the additional proposed constraint on *must*.

Suppose Holmes is hired as a consultant on the big case. The police can't afford his rates for solving the mystery entirely and so hire him to narrow things down to two suspects. Everyone knows that Holmes has his notebook of clues and knows that he never shares its contents: if his methods were disclosed, he'd soon be out of work.

Footnote 34 continued

- (i) (Alex is heading outside with no umbrella, galoshes, or raingear.)

Billy: You must not know that it is raining out.

Billy: You must not realize that it is raining out.

Billy is not trying to make it common ground between Alex and Billy that Alex doesn't know that it is raining. In fact, the opposite. So drawing attention to a mutually salient and available argument in support of it would be self-defeating. Perhaps, as suggested to us by Angelika Kratzer, the official proposal can be amended by carefully balancing the time at which Alex doesn't know that it is raining and the time at which the common ground is updated. Maybe so, but it seems tricky.

(32) (Holmes consults his notebook, puts it back in his breast pocket, and clears his throat.)

Holmes: The gardener can't be the murderer. It must be the butler or the driver.

Both modal claims are fine here even though Holmes doesn't and won't reveal an argument for *It isn't the gardener* and *It is the butler or the driver*. But the signal of indirectness remains.³⁵

We conclude that *must*'s evidential distribution is both too flexible and too persistent to be given this sort of pragmatic explanation.

5.3 Manner

Mandelkern's explanation is a sort of hybrid: he argues that the novel pragmatic constraint about mutually salient arguments can be given a manner-based implicature derivation and then that in turn can underwrite *must*'s signal of indirectness. We have objected that the novel constraint isn't as widespread as the signal of indirectness and so can't ground that signal.

But maybe we are wrong about *must*'s evidential signal. Maybe it's not about indirectness at all (it is) but about some other thing that can be given a manner-implicature explanation. Then the broad outline of an explanation would go like this: *must* ϕ and ψ , in the relevant context, have the same semantic upshot. Of these, *must* ϕ is syntactically more complex. Why use the more complex expression to achieved the same net effect? To convey non-semantic extra content.

What is needed to take this from a just-so story to an explanation is non-trivial. We need a candidate for ψ , we need to know about syntactic complexity, and we need to know how and why *must* ϕ carries this extra, non-semantic information (and what it is). And we need to know all of this in a way that is: (i) non-magical; and (ii) general enough to cover the evidential distribution of both *must* and *can't*.

This checklist is daunting. We stand ready to consider such a worked-out S1+S2+W3 package deal, when it gets worked out. Until then, S3 will do.³⁶

³⁵ But wait, just because Holmes uses the contents of the notebook to rule out the gardener and narrow in on the butler and the driver, can't his *can't* and *must* point to a different argument? For instance, one that is constructable-on-the-fly to his audience that relies on the fact that Holmes consulted his notebook and used *must*? No. First, because the proposed pragmatic derivation requires that the argument is the speaker's best evidence and not one based on the hearers trusting that the speaker has some private evidence for the prejacent. Second, because such on-the-fly arguments are too easily constructed, threatening to predict the acceptability of *must* in situations where it isn't—including, for instance, in reporting Patch's antics in (31b). (Thanks again to Angelika Kratzer for input on this.)

³⁶ We note that Swanson (2008) (see esp. fn.14) argues that hardwiring the evidential signal may not be so bad after all. We're not sure we agree with his particular reasons, but hey, we appreciate the support. Finally, we note that Matthewson and Truckenbrodt (2018) show subtle differences between English *must* and German *müssen*, which according to them argue in favor of a semantic hardwiring of the evidential requirement.

6 Conclusion

We conclude that *must* must be strong. We're maximally confident of that.

References

- Del Pinal, Guillermo, and Brandon Waldon. 2019. Modals under epistemic tension. *Natural Language Semantics* 27 (2): 135–188. <https://doi.org/10.1007/s11050-019-09151-w>.
- von Fintel, Kai. 2001. Counterfactuals in a dynamic context. In *Ken Hale: A life in language*, ed. Michael Kenstowicz, 123–152. Cambridge, MA: MIT Press.
- von Fintel, Kai, and Anthony S. Gillies. 2007. An opinionated guide to epistemic modality. In *Oxford studies in epistemology*, vol. 2, ed. Tamar Szabó Gendle and John Hawthorne, 32–62. Oxford: Oxford University Press.
- von Fintel, Kai, and Anthony S. Gillies. 2008. CIA leaks. *The Philosophical Review* 117 (1): 77–98. <https://doi.org/10.1215/00318108-2007-025>.
- von Fintel, Kai, and Anthony S. Gillies. 2010. *Must...stay...strong!* *Natural Language Semantics* 18: 351–383. <https://doi.org/10.1007/s11050-010-9058-2>.
- von Fintel, Kai, and Anthony S. Gillies. 2011. 'Might' made right. In *Epistemic modality*, ed. Andy Ega and Brian Weatherston, 108–130. Oxford: Oxford University Press.
- Giannakidou, Anastasia. 1999. Affective dependencies. *Linguistics and Philosophy* 22 (4): 367–421. <https://doi.org/10.1023/A:1005492130684>.
- Giannakidou, Anastasia, and Alda Mari. 2018. A unified analysis of the future as epistemic modality. *Natural Language and Linguistic Theory* 36 (1): 85–129. <https://doi.org/10.1007/s11049-017-9366-z>.
- Gillies, Anthony S. 2007. Counterfactual scorekeeping. *Linguistics and Philosophy* 30 (3): 329–360. <https://doi.org/10.1007/s10988-007-9018-6>.
- Goodhue, Daniel. 2017. *Must ϕ* is felicitous only if ϕ is not known. *Semantics & Pragmatics* 10 (14): 1–27. <https://doi.org/10.3765/sp.10.14>.
- Grice, Paul. 1967. Logic and conversation. Lecture Notes for William James lectures at Harvard University, published in slightly revised form in H.P. Grice, *Studies in the way of words*, pp. 1–43, Harvard University Press, 1989. 1–143.
- Groenendijk, Jeroen A.G., and Martin J.B. Stokhof. 1975. Modality and conversational information. *Theoretical Linguistics* 2 (1/2): 61–112. <https://doi.org/10.1515/thli.1975.2.1-3.61>.
- Ippolito, Michela. 2018. Constraints on the embeddability of epistemic modals. *Sinn und Bedeutung (SuB)* 21: 605–622.
- Jacobson, Pauline. 2018. What is—or, for that matter isn't—'experimental' semantics? In *The science of meaning: Essays on the metatheory of natural language semantics*, 46–72, ed. Derek Ball and Brian Raburn. Oxford: Oxford University Press.
- Karttunen, Lauri. 1972. Possible and must. In *Syntax and semantics*, vol. 1, ed. John P. Kimball, 1–20. New York: Academic Press.
- Kratzer, Angelika. 1991. Modality. In *Semantics: An international handbook of contemporary research*, ed. Arnim von Stechow and Dieter Wunderlich, 639–650. Berlin: de Gruyter.
- Križ, Manuel. 2015. Homogeneity, non-maximality, and *all*. *Journal of Semantics*. <https://doi.org/10.1093/jos/ffv006>.
- Kroch, Anthony. 1974. The semantics of scope in English. PhD thesis, MIT. 1721.1/13020.
- Lasnik, Peter. 1999. Pragmatic halos. *Language* 75 (3): 522–551. <https://doi.org/10.2307/417059>.
- Lassiter, Daniel. 2016. *Must*, knowledge, (in)directness. *Natural Language Semantics* 24: 117–163. <https://doi.org/10.1007/s11050-016-9121-8>.
- Lewis, David. 1979. Counterfactual dependence and time's arrow. *Noûs* 13 (4): 455–476. <https://doi.org/10.2307/2215339>.
- Lyons, John. 1977. *Semantics*. Cambridge: Cambridge University Press.
- Mandelkern, Matthew. 2016. A solution to Karttunen's problem. *Sinn und Bedeutung* 21. <http://users.ox.ac.uk/sfop0776/KarttSuB.pdf>.
- Mandelkern, Matthew. 2019. What *must* adds. *Linguistics and Philosophy* 42 (3): 225–266. <https://doi.org/10.1007/s10988-018-9246-y>.

- Matthewson, Lisa, and Hubert Truckenbrodt. 2018. Semantic building blocks: Evidential epistemic modals in English, German and St'át'imcets. In *Thoughts on mind and grammar: A festschrift in honor of Tom Roeper* (UMOPL 41), ed. Bart Hollebrandse, Jaieun Kim, Ana T. Pérez-Leroux, and Petra Schulz, 89–100. Amherst, MA: GLSA.
- Ricciardi, Giuseppe, Rachel Ryskin, and Edward Gibson. 2020. Epistemic *must p* is literally a strong statement. Poster presented at CUNY Sentence Processing Conference. <https://osf.io/rwpag>.
- Sherman, Brett. 2018. Open questions and epistemic necessity. *The Philosophical Quarterly* 68 (273): 819–840. <https://doi.org/10.1093/pq/pqy025>.
- Swanson, Eric. 2008. Modality in language. *Philosophy Compass* 3 (6): 1193–1207. <https://doi.org/10.1111/j.1747-9991.2008.00177.x>.
- Veltman, Frank. 1985. Logics for conditionals. PhD thesis, University of Amsterdam. <https://www.illc.uva.nl/cms/Research/Publications/Dissertations/HDS-02-Frank-Veltman.text.pdf>.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.