

Epistemic *must* and *might*: evidence that argumentation is semantically encoded

Brandon Waldon
Stanford University

1 What this paper is about

This paper is about epistemic *must* and *might*. A sizeable proportion of the literature on epistemic modality takes as its point of departure data such as the following:

- (1) (In the context of direct observation of rain):
 - a. # It must be raining outside.
 - b. It is raining outside.

Karttunen (1972) remarks that *must* ϕ is degraded when compared to bare unmodalized assertions of ϕ in ‘direct-observation’ contexts like (1). I’ll briefly address this ‘directness’ problem for the semantics of *must*, but my point of departure is a less-discussed issue. I’ll call it the ‘argumentation’ problem, discussed at length by Stone (1994) and illustrated by exchanges such as the following:

- (2) a. A: *What’s the weather like outside?*
b. B: ? *It must be raining.*

Stone remarks that in order for (2b) to be felicitous, Speaker A has to be able to identify the particular chain of reasoning that licenses the conclusion that it’s raining. If A cannot reconstruct B’s argument for *rain*, then (2b) is degraded.

Like the directness problem, the argumentation problem has to do with an asymmetry in the felicity conditions of *must* ϕ and its bare counterpart. In (2), if B simply responds *It’s raining*, then B’s argument for *rain* need not be salient in context.

A solution to the argumentation problem explains this asymmetry. I review two perspectives on the problem: the Semantic account, whereby the ‘argumentation’ property of *must* is encoded in the semantics of the modal; and the Pragmatic account, which posits that this property is contextually conditioned. Next, I present experimental evidence against the Pragmatic account. Finally, I present a solution to the argumentation problem and offer an account of *must*’s putative dual, *might*.¹

2 Two perspectives on the argumentation problem

2.1 The Semantic account (Stone 1994)

Stone (1994)’s Semantic solution to the argumentation problem is to encode a notion of ‘argumentation’ directly in the semantics of *must*. The analysis I ultimately propose in Section 5 extends Stone’s argument system semantics for *must* and bare assertions, the important aspects of which I provide below:

¹I’d like to thank Christian Brickhouse, Cleo Condoravdi, Judith Degen, Atticus Geiger, Dan Lassiter, Chris Potts, and three CLS reviewers for their insightful feedback on this work.

(3) $\underline{\phi}$ is true in \mathcal{K} iff $\mathcal{K} \models \phi$ (adapted from Stone 1994: p. 7)

(4) $\underline{\text{Must } \phi}$ is true in \mathcal{K} iff $\mathcal{K} \models \langle A, \phi \rangle_{\mathcal{K}}$ (adapted from Stone 1994: p. 8)

... where \mathcal{K} is a conversational context whose formal type is an argument system (a set of ‘established propositions’ as well as both logical and defeasible rules of inference) and A is an argument: a contextually salient, non-empty subset of the defeasible rules in \mathcal{K} . In the case of (3), $\mathcal{K} \models \phi$ simply denotes that \mathcal{K} justifies concluding ϕ . In (4), $\mathcal{K} \models \langle A, \phi \rangle_{\mathcal{K}}$ denotes that concluding ϕ on the basis of A is justified in \mathcal{K} .²

When we compare (3) and (4), we find a straightforward solution to the argumentation problem. In the case of *must* ϕ , the listener must be able to identify the particular intended salient collection of premises A in the context; otherwise, we expect a crash. No such requirement exists for bare assertions, whose truth in an argument system simply makes reference to the argument system as a whole.³

Stone’s Semantic account also provides a straightforward solution to the directness problem, if we make additional assumptions about what constitutes a well-formed argument and about the nature of the context. On Stone’s analysis, an argument A for a conclusion ϕ in \mathcal{K} must consist of the minimal set of defeasible rules of inference which - coupled with the set of established propositions and the logical rules of inference in \mathcal{K} - entails ϕ . If ϕ is already in the set of established propositions, then the minimal set of required premises is empty. The prediction is that there is no non-empty argument A that can meet the definedness conditions of *must* ϕ if ϕ is already established in \mathcal{K} . If we furthermore assume that direct visual evidence of ϕ adds ϕ to \mathcal{K} automatically, then we predict that *must* ϕ should be infelicitous in ‘direct-observation’ contexts such as (1).⁴

2.2 The Pragmatic account (Mandelkern 2019)

Mandelkern (2019), following Stone (1994), notes that *must* ϕ requires that the listener can identify the speaker’s argument for concluding ϕ . Contrasts such as (5b) and (5c) provide further illustration of the argumentation problem, while (5d) demonstrates that the speaker can satisfy the felicity conditions on *must* ϕ by explicitly providing evidence for ϕ (data from Mandelkern 2019: 251):

- (5) (Context: Julie’s cat has been sneezing a lot. Ben asks her how the cat is doing. Julie says:)
- a. Not so great. I need to take him to the vet actually, he might have an upper respiratory infection.
 - b. Not so great. I need to take him to the vet actually, he has an upper respiratory infection.

²Stone’s semantics builds on the argument system approach to defeasible reasoning found in Simari & Loui (1992). Roughly, an argument for ϕ is ‘justified’ in an argument system if the argument - coupled with other premises in the system - entails ϕ and there is no argument that makes use of more information in the system to generate the conclusion $\neg\phi$.

³Stone’s account is officially silent with regards to how assertions of (*must*) ϕ can lead to an update of the context. I will also have to punt on this issue for the time being.

⁴von Stechow & Gillies (2010) offer a similar solution to the directness problem: on their account, *must* ϕ is undefined if ϕ is already ‘settled’ in context by virtue of ϕ ’s direct observation.

- c. ?? Not so great. I need to take him to the vet actually, he must have an upper respiratory infection.
- d. Not so great. I need to take him to the vet actually, he must have an upper respiratory infection; he's been sneezing a lot lately.

However, Mandelkern also notes that the acceptability of (5a) - which features a *might* ϕ claim but no contextually salient argument for ϕ - poses a challenge for the Semantic account. If *might* and *must* are duals, as is commonly assumed, then they should have similar definedness conditions.

Mandelkern argues that all epistemic modals highlight interlocutors' shared evidence for the proffered conclusion (sometimes called the 'prejacent' of the modal) as part of their conventional meaning, but *must* ϕ is unique in that it competes pragmatically with an informationally-equivalent alternative, namely the bare form ϕ . On Mandelkern's analysis, a speaker's choice to produce *must* ϕ sets off a chain of reasoning for the listener, schematized in (6) below. The salient-argument felicity condition on *must* is cached out as a manner inference, expressed in (6h).⁵

- (6) a. A hears B assert *Must* ϕ .
- b. *Must* ϕ has a structurally simpler and informationally equivalent alternative, namely ϕ .
- c. So A will ask why B chose to assert *Must* ϕ rather than ϕ in order to propose an update with ϕ .
- d. The only relevant difference between the two is that *Must* ϕ , and not ϕ , calls attention to the interlocutors' common evidence for ϕ .
- e. So A concludes that B finds this evidence noteworthy.
- f. If the interlocutors' common evidence for ϕ were just that B had proposed to update with ϕ , and thus endorses ϕ , there would be nothing noteworthy about this common evidence, since whenever a speaker proposes to update with ϕ , she thereby indicates that she endorses ϕ .
- g. So A concludes that B intends to update with ϕ on the basis of some shared evidence that goes beyond the simple fact that B endorses ϕ ; in other words, on the basis of some substantive shared evidence in support of ϕ .
- h. Thus in order for B's assertion of *Must* ϕ to be felicitous, B must ensure that there is a salient argument which she endorses as an argument for ϕ .
(Mandelkern 2019: 254-255)

Because sentences embedded in weaker epistemic modals (including *might* and *should*) are not informationally equivalent to unmodalized sentences (that is, there is no analogous premise 6b for the weaker modals), Mandelkern's analysis predicts that the salient-argument felicity condition does not apply to weaker expressions.⁶

⁵A similar manner-inference proposal is made by Degen *et al.* (2015), who argue that the choice to produce the marked *must* ϕ over the simpler bare ϕ signals to the listener that the speaker's evidence for ϕ is likely weak.

⁶I won't discuss *should*, but Stone and Mandelkern each note that the argumentation problem doesn't seem to apply to it. Stone provides a semantics for *should* that resembles the analysis of

Furthermore, on Mandelkern's full account, the salient argument highlighted by *must* is subject to additional restrictions. In particular, it has to be both 'non-redundant' and the strongest available argument in support of the conclusion.⁷ This provides a solution to the directness problem: the listener who encounters (1a) concludes - via (6) - that the speaker intends to highlight a particular argument for 'it's raining outside.' However, the strongest salient argument for 'it's raining outside' in (1) is a redundant one ('it's raining outside because we can see it's raining outside'); therefore, (1a) is infelicitous.

2.3 Predictions

The Semantic account locates the felicity condition expressed by (6h) in the semantics of the modal itself. Therefore, the Semantic account predicts that this felicity condition should be context-independent.

Conversely, the Pragmatic account predicts that the salient-argument felicity condition on *must* should be context-dependent; that is, it should only emerge in contexts where listeners recognize that the speaker had a simpler alternative at her disposal (the bare assertion) that would have made an informationally-equivalent proposal to update the common ground.⁸ This account therefore predicts that there should exist no context *C* such that:

1. *Must* ϕ is assertible in *C*;
2. The differences between bare ϕ and *must* ϕ in *C* go beyond 'calling attention to common evidence'; and
3. The salient-argument felicity condition on *must* persists in *C*.

If the second clause holds, then we deny premise (6d) above. But on the Pragmatic account, (6d) is a necessary step to (6h), which expresses the salient-argument felicity condition on *must*. So the second and third clauses above cannot both be true simultaneously on the Pragmatic account. In the following section, I identify a class of contexts where these conditions do appear to hold simultaneously.

3 A potential wrinkle in the Pragmatic account: subjective predicates in the scope of *must*

Malamud & Stephenson (2014) discuss the conventional effects of reverse-polarity tag questions (as in the tag question in 7a), same-polarity tag questions (7b), and rising intonation (7d) in so-called 'blushing/innuendo' contexts such as (7):

must offered by Kratzer (1991), whereby *should* ϕ marks that ϕ follows what is known at the world of evaluation, coupled with assumptions about what is true of a 'normal' world. Mandelkern takes the behavior of *should* to be further support for his Pragmatic account. But see also Copley (2006) and Yalcin (2016), who provide evidence that 'epistemic' *should* is not truly epistemic in flavor.

⁷For Mandelkern, non-redundancy is a general pragmatic prohibition against "proposal(s) to update the common ground with *p* on the basis of an argument Γ ... if *p* follows from Γ in a way that is mutually recognized to be obvious" (2019: 242).

⁸A lot of ink has been spilled over the question of what speakers commit to with *must* ϕ . If listeners consider *must* ϕ to make a weaker proposal than does bare ϕ , then premise (6b) above is in question. However, I won't discuss this issue here. (For recent experimental investigations see e.g. Lassiter 2016; Degen *et al.* 2019; Del Pinal & Waldon 2019; Ricciardi *et al.* 2020).

- (7) (Context: A and B are gossiping. A doesn't know anything about B's neighbor. B says, blushing, 'You've GOT to see this picture of my new neighbor!' **Without looking**, A replies:)
- a. # He's attractive, isn't he?
 - b. He's attractive, is he?
 - c. He's attractive?
 - d. # He's attractive. (Malamud & Stephenson 2014: 279)

Malamud & Stephenson remark that in (7), "B's judgment of attractiveness is at issue and A's is not. Therefore a felicitous effect of A's move must have something to do with B's commitments to the anchor proposition ['the neighbor is attractive']" (2014: 279). The infelicity of (7a) and (7d), then, stems from an incongruity between what is at-issue in context (B's judgment) and what is made at-issue via A's assertion. Both (7a) and (7d) are most straightforwardly interpreted as expressing A's thoughts about B's neighbor, but as the context makes clear, A is in no position to make a judgment about B's neighbor's attractiveness.

However, in the scope of *must*, the taste predicate *attractive* is assertible by A, as in (8a) (8b repeats 7d to illustrate the contrast with the bare form):

- (8) (In the context of 7):
- a. He must be attractive.
 - b. # He's attractive.

Native speakers I've consulted tend to interpret (8a) as a comment on B's judgment of the neighbor's attractiveness, something akin to *it must be the case that B finds the neighbor attractive*. This is in line with Malamud & Stephenson's assessment of what's at issue in the context of (7). Another available interpretation seems to be: *it must be the case that I (Speaker A) would find the neighbor attractive (given a chance to see the neighbor)*.⁹

Regardless of how one interprets (8a), it's clear that the proposal made by the assertion is qualitatively distinct from the one made by asserting (8b), which strongly favors the reading *I (Speaker A) find the neighbor attractive*.¹⁰

⁹Stephenson (2007) similarly observes that when a modal scopes over a subjective predicate - as in *The cat food might be tasty* - the epistemic anchor of *might* can only be the speaker; however, the sentence has two available readings whereby the speaker or some third party (e.g. the cat eating the cat food) finds the cat food tasty.

¹⁰Lasersohn (2005) presents examples where subjective predicates in unmodalized sentences can apparently receive 'exocentric' readings, where the relevant experiencer is not the speaker: (Context: John is describing to Mary how their two-year-old son Bill enjoyed a recent trip to the amusement park):

- a. Mary: *How did Bill like the rides?*
- b. John: *Well, the merry-go-round was fun, but the water slide was a little too scary.*

What is important for my purposes is that - for whatever reason - contexts such as (7) are ones where the non speaker-oriented reading of, e.g., *attractive*, is evidently much less available in the scope of a bare sentence than in the scope of *must*.

So contextual premise (6d) does not hold in (8); there are multiple relevant differences between *must* ϕ and bare ϕ beyond the extent to which each form calls attention to ‘common evidence.’ Thus, when B hears (8a), she can explain away A’s decision not to produce a structurally simpler utterance - (8b) - on grounds that have nothing to do with the notion of ‘common evidence.’

However, if (6d) does not hold as a contextual premise, then the Pragmatic account predicts that the inference to (6h) should not go through and that *must* ϕ should be felicitous without a contextually salient argument for ϕ . This seems dubious in the context of (8): the felicity of (8a) seems to hinge on the recognition that one can reason to the conclusion from an argument that hinges on salient contextual evidence (i.e. the fact that B is blushing; the fact that B is so excited to show A the picture of the new neighbor). Without this evidence, (8a) is degraded.

(8) appears to be a context that the Pragmatic account predicts should not exist: *must* ϕ is assertible, and ϕ makes a distinct proposal from the one made by *must* ϕ ; nonetheless, the salient-argument felicity condition on *must* ϕ emerges.

In the following section, I report the results of an experiment designed to systematically investigate the acceptability of subjective predicates in contexts such as (8). The common feature of the contexts of interest is that, as in (8), the speaker does not have the requisite direct sensory experience that would allow her to assert some subjective predicate (e.g. *attractive*) in the scope of an unmodalized sentence. Nonetheless, the speaker is apparently able to assert that same predicate in the scope of *must*. I will call these contexts “secondary-experience” contexts in what follows.

4 Experiment

4.1 Methods

225 participants were recruited using Amazon Mechanical Turk (native English speakers; US IP Addresses; 95% prior approval rating on MTurk)¹¹.

Participants were instructed to “read a short passage in which two individuals are interacting with one another. At the end of the passage, one of the individuals says a sentence (in **bold**).” The bolded sentence could appear in one of three conditions: a *must*-condition, where a subjective predicate occurs in the scope of *must*; a ‘bare’ condition, where the predicate occurs in an unmodalized sentence; and a ‘good’ control condition, where the bolded sentence merely repeats material made explicit in the passage. One such vignette is presented below:

(9) ‘Exam’ vignette:

Brian meets Jill shortly after she finishes her final calculus exam. Brian didn’t take the exam but can see that Jill is visibly distressed.

Brian says: [It must have been a difficult exam! (must) / It was a difficult exam! (bare) / Jill, you look distressed! (control)]

There were three conditions (*must* ϕ , bare ϕ , and control) fully crossed with five vignettes for a total of 15 critical conditions. The experiment featured a fully

¹¹This experiment was preregistered at <https://osf.io/f6snq>. Data and code are accessible from <https://github.com/bwaldon/probmust>.

between-subjects design. I collected an average of 15 responses per condition.

4.1.1 Procedure

Participants were told to rate the extent to which the speaker "was in a position to say [the bolded] sentence, given what you've read in the short passage".¹² Participants read the vignette and were provided a 7-point scale: the negative endpoint 1 was labeled "Not at all in a position to say the **bolded** sentence"; the positive endpoint 7 was labeled "Definitely in a position to say the **bolded** sentence".

After they provided their rating on the 7-point scale, participants answered a short comprehension question as an attention check. These questions consisted of a forced choice between two options and tested the participants' knowledge about a basic fact of the story (for example, for the vignette in 9: *What did Jill just do?* with a choice between 'Write an exam' and 'Run a marathon').

4.2 Experiment predictions

In the control condition, the speaker's bolded sentence merely repeats information provided in the context. Therefore, I expect participants to rate the speaker as being "Definitely in a position" to assert the bolded material.

Regarding the bare condition, I assume (following Malamud & Stephenson 2014) that "when X asserts or otherwise presents themselves as believing, e.g. that Y is *attractive*, this typically conveys that Y is attractive as judged by X, but not necessarily that Y is attractive as judged by other participants in the conversation" (2014: 277); that is, a subjective predicate in the root clause of an unmodalized sentence typically conveys that the speaker is making a subjective evaluation relative to her own perspective. Because the speaker does not have the requisite direct sensory experience that would allow her to make such an evaluation in secondary-experience contexts, I expect ratings in the bare condition to fall towards the lower end of the scale.

However, the secondary-experience contexts in this experiment signal that there is adequate evidence to associate the addressee (if not the speaker) with a subjective evaluation (for example, in 9: Jill's visible distress indicates that *the exam was difficult for Jill*). If the *must*-condition sentences are assertible on this addressee-oriented meaning, then I expect the *must*-sentences to pattern with the controls. This is because I expect that participants will judge the speakers to be in a position to make addressee-oriented claims given the evidence provided in the context.

By similar reasoning, I also expect the *must* condition sentences to receive high ratings if they receive the second available interpretation (that is, if the sentences are interpreted as statements about the speaker's own subjective commitments conditional on direct sensory experience - for example, in 9: *I (Brian) would have found the exam difficult had I taken it*). Thus, regardless of which of the two interpretations participants actually draw in the *must* condition, I predict that ratings in that condition will be higher than ratings in the bare condition.

¹²Instructions continue: "For example, if you know from the passage that Bob is standing outside and can see that it's sunny, then he's in a position to say **It is sunny outside**. But if he's sitting inside (and has no reason to believe it's sunny outside), then he's not in a position to say that sentence."

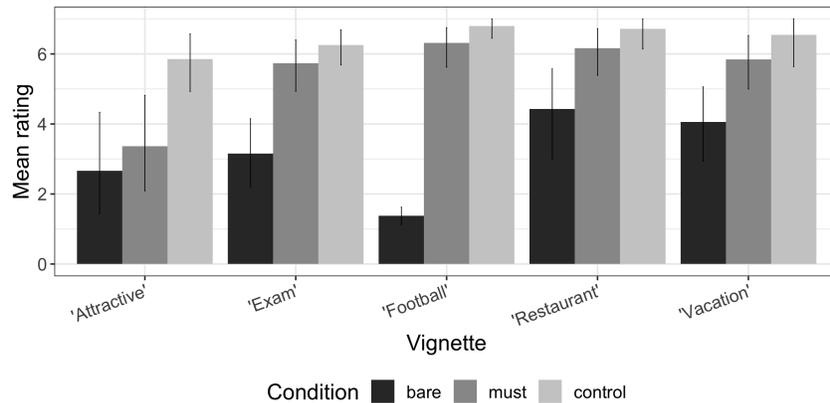


Figure 1: Experiment results, subsetted by vignette. Error bars indicate 95% bootstrapped confidence intervals.

4.3 Results

Four responses were excluded because they came from participants who reported a native language other than English. A further five responses were removed due to participants failing the attention check. In total, I excluded 4% of the data.

Figure 1 displays the results of the experiment. As predicted, behavior in the control condition is basically at ceiling across all five vignettes; moreover, for each vignette, the mean ratings in the *must* condition (which are mostly clustered towards the high end of the scale) are higher than the ratings in the bare condition.

To assess whether there was evidence of a contrast in ratings between the *must* and bare conditions, a pairwise Bayesian ordinal logistic regression model was fit to the data subset to those two conditions, predicting response from a fixed effect of condition (with the maximal random effects structure justified by the design, i.e. random by-vignette slopes for condition, and the bare condition as the reference level). The model suggested strong evidence of an effect of condition (95% credible interval of coefficient: [0.66, 3.67]).¹³

The comparison of ratings between the *must* and bare conditions is informative, but it is also instructive to consider the absolute positions of the mean ratings in the bare condition across the five vignettes. In each case, the mean rating is at or below the middle of the scale. I cautiously interpret this result as suggesting that in the bare condition, participants were overall uncertain (at best) as to whether the speaker was in a position to make the assertion. Notably, the worst-performing *must*-condition vignette was the vignette that minimally extended the original ‘blushing/innuendo’ example from Malamud & Stephenson (2014):

(10) ‘Attractive’ vignette:

Annie and Betty are gossiping. Annie doesn’t know anything about Betty’s neighbor. Betty says, blushing, “You’ve GOT to see this picture of my new neighbor!”

Without looking at the picture, Annie replies: [He must be attractive! (must) / He’s attractive! (bare) / Betty, you’re blushing! (control)]

¹³I implemented the regression with the `brms` package (Bürkner 2018) in R, using the default priors over model parameters supplied by `brms`.

It's possible that the evidential cue in this context (i.e. Betty's blushing) was not prominent enough to signal to most participants that Annie was in a position to make a claim about Betty's subjective evaluation of the neighbor's attractiveness (or that she was in a position to make a claim about her own conditional subjective evaluation of the neighbor, depending on how *must be attractive* is interpreted by the participant). Note that *blushing* was a single-word cue in (10) and it occurred in an adjunct of one of the sentences in the passage; in these two respects, it differs from the cue in the 'Exam' vignette in (9). If participants did not pick up the cue in (10), then we would expect the *must*-condition sentence to be degraded.

4.4 Discussion

The results of the experiment support the claim that *must*'s salient-argument felicity condition arises even in contexts where listeners cannot be plausibly engaged in the reasoning schematized by Mandelkern in (6). This poses an issue for any pragmatic analysis that derives the salient-argument felicity condition as a manner inference.

But can the Pragmatic account be altered to accommodate the pattern of results presented here? One possibility is that the salient-argument felicity condition arises due to *must* ϕ competing with some alternative other than the bare form. It is difficult to imagine what this alternative form would be: presumably, it should be a form of equal or lesser structural complexity than *must* ϕ (Katzir 2007), whose non-production can be explained exclusively by the speaker's desire to highlight 'common evidence' for the conclusion. I showed that in secondary-experience contexts, the speaker's decision not to produce the bare form can be explained on independent grounds, namely that the bare assertion would make a qualitatively distinct proposal on the common ground. I suspect that the same would be true of any other form we might consider to be the relevant alternative to *must* ϕ .

With no clear way to assimilate the Pragmatic account to my experimental data, I propose an analysis that extends the Semantic account offered by Stone (1994).

5 Analysis

The empirical picture presented thus far suggests the following desiderata of a solution to the argumentation problem:

- (11) Analysis desiderata (strong formulations, to be weakened):
1. **Desideratum 1:** The analysis should predict that *must*'s salient-argument felicity condition is context-independent;
 2. **Desideratum 2:** The analysis should account for why this felicity condition does not arise for weaker epistemic modals.

Stone's analysis meets this first desideratum but does not address the second. Though Stone does not offer a semantics for *might*, Mandelkern correctly points out that Stone's analysis - when extended to *might* - predicts that *must* and *might* should be subject to the salient-argument felicity condition to the same extent.

For a unified picture of *must* and *might*, I extend Stone's analysis, beginning with a revision of Stone's formal representation of conversational contexts. On

Stone’s original analysis, contexts are represented as individual argument systems - sets of atomic propositions and inferential rules. In contrast, I assume that the context does not definitively settle which argument system is the ‘correct’ one for the purposes of inference and decision making. That is, there may be uncertainty as to the truth/falsity of certain propositions, or there may be uncertainty as to whether certain defeasible rules of inference may be relied upon.

On my proposal, then, contexts are represented as sets of possible argument systems. For a proposition to be supported by a context, it must be supported by all of the constituent argument systems of that context.

(12) A context Z supports ϕ iff $\forall \mathcal{K}_{\in Z} : \mathcal{K} \models \phi$

On the strong quantificational variant of my analysis, *must* and *might* are universal and existential quantifiers over Z , respectively.

(13) Entries for *must* and *might*:¹⁴

a. Must ϕ is true in context Z iff $\forall \mathcal{K}_{\in Z} : \mathcal{K} \models \langle A, \phi \rangle_{\mathcal{K}}$

b. Might ϕ is true in context Z iff $\exists \mathcal{K}_{\in Z} : \mathcal{K} \models \langle A, \phi \rangle_{\mathcal{K}}$ ¹⁵

... where $\mathcal{K} \models \langle A, \phi \rangle_{\mathcal{K}}$ denotes that concluding ϕ on the basis of a contextually-salient collection of defeasible rules A is justified in \mathcal{K} .

I’ll illustrate with an example. Suppose speaker S is standing inside her house and notices that her roof is leaking. Let `rain` stand in for the proposition expressed by *It is raining outside S’s house*, and let `leakyroof` stand in for *S’s roof is leaking*. We can model S ’s representation of the context as a set of argument systems, all of which contain `leakyroof` (that is, we assume S is maximally committed to the proposition that her roof is leaking):

(14) Argument systems in S ’s representation of the context:¹⁶

$\mathcal{K}_1 : [\dots, \text{leakyroof}, \neg \text{hose}, \Delta]$ $\mathcal{K}_2 : [\dots, \text{leakyroof}, \neg \text{hose}, \Delta]$ $\mathcal{K}_3 : [\dots, \text{leakyroof}, \neg \text{hose}, \Delta]$ $\mathcal{K}_4 : [\dots, \text{leakyroof}, \text{hose}, \Delta]$ $\Delta = \{\text{leakyroof} \rightarrow \text{rain}, \text{leakyroof} \wedge \text{hose} \rightarrow \neg \text{rain}\}$
--

¹⁴On a weak variant of (13), *must* ϕ and *might* ϕ are true if the likelihood that an argument system in the context supports ϕ (on the basis of a salient argument) exceeds a contextual threshold.

For expository purposes, I stick to the simpler denotations in (13) in this paper. In Waldon (in press), I explore how the weak analysis - coupled with a Bayesian probabilistic formalism of pragmatic inference - allows us to model how speaker productions of *must* ϕ give rise to inferences both about the speaker’s commitment to ϕ and about the speaker’s beliefs regarding the structure of the context (see also Swanson 2006, Lassiter 2016 for similar probabilistic proposals).

¹⁵On (13), *must* ϕ can be false in Z when the salient argument A doesn’t justify ϕ in all of Z ’s argument systems. Stone (1994: ex. 36) offers some evidence that *must* ϕ assertions can indeed be the target of linguistic denial when the salient argument for ϕ is unjustified, even when all interlocutors agree that $\neg\phi$ is not an epistemic possibility. But that would mean - consistent with my proposal - that *must* and *might* are not necessarily duals: $\neg \text{must}(\neg\phi) \neq \text{might}(\phi)$.

¹⁶In (14), “. . .” stands for information that is irrelevant to whether or not `rain`.

Let `hose` stand for the proposition expressed by *S's neighbor is spraying S's roof with a hose*. The possible argument systems contained in (14) differ as to whether or not `hose` or its negation holds. That is, *S* is uncertain as to whether or not her neighbor is outside spraying the roof with a hose (though more likely than not, this is not the case).

Each argument system in (14) contains the same set of defeasible inferential rules. If the roof is leaking, then it is raining outside; if the roof is leaking but it is additionally revealed that the neighbor is spraying *S*'s roof with a hose (an unlikely occurrence), then it is not raining outside.

\mathcal{K}_1 , \mathcal{K}_2 , and \mathcal{K}_3 justify the conclusion `rain` on the basis of the argument `leakyroof` \rightarrow `rain`; however, \mathcal{K}_4 does not: another argument in the system makes use of more information to generate a conclusion that contradicts `rain`.

Given this setup (and given 13), *might*(`rain`) is true of (14), provided the argument made salient in context is `leakyroof` \rightarrow `rain`. So is *might*(\neg `rain`), provided the salient argument is `leakyroof` \wedge `hose` \rightarrow \neg `rain`. However, *must*(`rain`) is false if the salient argument is `leakyroof` \rightarrow `rain`, because that argument for `rain` is not justified in all of the possible argument systems.

The analysis suggests that *must* is undefined if there is no argument *A* made salient in the context, so we satisfy Desideratum 1. However, the same is suggested for *might*, so more must be said in order to fully address Desideratum 2. I start by recapitulating Stone's remarks on data such as (15) (from Stone 1994: 4):

- (15) a. Ann: *Where is the sugar?*
 b. Mary: *It must be in the cabinet over the fridge.*

“[T]o make sense of Mary's answer in [15], **one must assume that Mary has just seen something or figured something out** from which she concludes that the sugar is in the cabinet over the fridge. Perhaps Mary has seen a telltale trail of white particles, or perhaps she has realized that only one cabinet remains in the kitchen which Ann has not ruled out” (Stone 1994: 4, emphasis added).

Stone's analysis suggests an accommodation strategy in case the listener observes *must/might* ϕ without a salient argument for ϕ : simply accept that the argument for ϕ exists and register the speaker's degree of commitment to ϕ . That is, when the context doesn't identify a salient argument, listeners accommodate the following presupposition, which I append to the entries in (13):

- (16) The argument presupposition of *must/might*:
 There exists a non-empty argument *A* for ϕ .¹⁷

We can assume without issue that this presupposition projects, e.g. in the scope of questions or under negation: even if *must/might* ϕ presupposes that an argument

¹⁷But for the same reason that anaphoric pronouns with no identifiable antecedents are hard to accommodate, it should be hard to accommodate *must/might* without a contextually supplied value for *A*: *must/might* ϕ are undefined on my semantics without such a valuation. This is a weakness of my current proposal, and I'm grateful to Cleo Condoravdi (p.c.) for identifying it.

A exists for ϕ , the speaker is not committed by virtue of that presupposition to anything regarding that argument's justifiability. Now the question is: why is accommodating this presupposition evidently less natural with *must* than with *might*?

To answer this, we first have to acknowledge that both Desiderata 1 and 2 oversimplify the empirical picture. First, it is evidently not the case that the salient-argument felicity condition arises in all contexts for *must*. von Fintel & Gillies (in press) identify contexts such as (17), where *must* is assertible even though the speaker does not indicate - implicitly or explicitly - an argument for the conclusion:

- (17) (Context: Holmes is hired as a consultant on a big murder case. The police can't afford his rates for solving the mystery 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 are disclosed, he'd soon be out of work. Holmes consults his notebook, puts it back in his breast pocket, and clears his throat.)

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

(adapted from von Fintel & Gillies in press)

Similarly, (18) appears to be a context where *must* can stand on its own:

- (18) (Context: Brian is upstairs, looking for his favorite shirt. He calls to his partner, Joe, who is watching TV downstairs:)
- a. Brian: Hey, where's my favorite shirt?
 - b. Joe: It must be in the washing machine.

Contexts such as (17) and (18) suggest that the unnaturalness of asserting *must* sans argument hinges on the relevance of the argument itself in context. In murder investigations, it's important not only to identify suspects but also to track the evidence and argumentation that implicates some individuals while exonerating others. That is, the context in (17) demands an answer to *Who is the murderer?* as well as - implicitly - to *Why do we have reason to (dis)believe that X is the murderer?*

In light of those demands, it would be odd for an investigator to declare *The gardener is the murderer!* without justifying the conclusion to her colleagues. But *The gardener must be the murderer!* is even worse, because *must* ϕ conventionally links ϕ to a set of inferential premises whose identities matter in context. Highlighting those premises via assertion of *must* - then failing to name them - is therefore especially uncooperative. An idiosyncratic exception, of course, is made for Holmes, whose narrow task is to address *Who is the murderer?*

The stakes are lower in (18): in uttering (18b), Joe indicates he has reason to think that shirt is in the washing machine. But Brian is only interested in answering *Where is my shirt?* - not *Why do we have reason to believe my shirt is in place X?* It suffices for Brian that Joe has an argument that justifies his certainty.

Additionally, there are contexts in which the acceptability of *might* ϕ is clearly bolstered by the presence of an argument in support of ϕ . Murder-mystery contexts such as (17) are one example. In (17), if Holmes is expected to share his clues

with investigators, then it's odd for him to say *The gardener might be the murderer!* without additional support for that claim.¹⁸

Desiderata 1 and 2, then, are generalizations of *must/might* that admit of exceptions. But paradigms such as (5) suggest that there is a true asymmetry between *must* and *might*: all else equal, the salient-argument felicity condition is stronger for the former than for the latter.

My analysis helps to address this asymmetry. Consider a Speaker 1 who asserts *must* ϕ on the basis of argument *A*. Speaker 1 thus conveys a high degree of certainty that *A* justifies concluding ϕ . Conversely, a Speaker 2 who asserts *might* ϕ on the basis of *A* in the same context (and who obeys Grice's Maxim of Quantity - to be sufficiently informative for the purposes of the conversational exchange) is relatively uncertain that *A* justifies ϕ . If we assume that high-certainty arguments are - all else equal - more contextually relevant than arguments that we consider unlikely to be justified, then we expect that Speaker 1 will feel under more pressure than Speaker 2 to make her argumentation clear to the listener.

5.1 A brief note on *certain* and *possible*

Like *must*, *certain* appears degraded in contexts of direct evidence for the prejacent:

- (19) (In the context of direct observation of rain):
- a. # It's certain that it's raining outside.
 - b. It is raining outside.

Moreover, it appears as though *certain* does not carry the salient-argument felicity condition to the same extent that *must* does. In the following examples from the Corpus of Contemporary American English (Davies 2009), *certain* ϕ is asserted without a contextually salient argument for ϕ . I replace *certain* with *must* to highlight the contrast in acceptability between these constructions in these contexts. With the *must* ϕ modifications (presented in brackets), the reader needs more explicit argumentation for ϕ compared to the original *certain* ϕ variants:

- (20) a. While it is possible each of us may be able to provide for ourselves and family - **it is certain** that we can not survive without the help of others.
[... ? - we **must** not be able to survive without the help of others.]
- b. ... the public has been kept in the dark about vaccination. **It is certain** that undisclosed, unlooked for illness occurs as a result of vaccines...
[... ? Vaccines **must** cause undisclosed, unlooked for illness ...]

¹⁸The salient-argument felicity condition can also arise for *might* if the conclusion proffered by the speaker is highly surprising. This is to be expected, if evidence in support of surprising claims is generally more contextually relevant than evidence in support of mundane ones:

(Context: Ann asks Mary: *Would you like to grab lunch on Friday?*)

- a. Mary: I'm not sure that works - I might have a work conflict that day.
- b. Mary: ??? I'm not sure that works - a meteor might hit Earth that day.
- c. Mary: I'm not sure that works - a meteor might hit Earth that day. Scientists have been tracking a large one headed to our solar system.

- c. Whether men in the Seventh Cavalry carried inverted flags during the Battle of the Little Bighorn is unknown, but **it is certain** that the Indians witnessed these soldiers in utmost distress on that day.
[... ? but the Indians **must** have witnessed these soldiers in utmost distress on that day.]

For *certain* ϕ to be true, we want each possible argument system in context to support ϕ (assuming that *certain* is at least as strong as *must*), but we do not want to require that there is a unique salient argument A that justifies the conclusion of ϕ in each of those states: this would predict felicity conditions for *certain* akin to those for *must*. The solution is to existentially bind the A variable such that it need not be valued by context:

- (21) Entries for *certain* and its putative dual, *possible*:
- a. Certain ϕ is true in context Z iff $\exists A : \forall \mathcal{K}_{\in Z} : \mathcal{K} \models \langle A, \phi \rangle_{\mathcal{K}}$
 - b. Possible ϕ is true in context Z iff $\exists A : \exists \mathcal{K}_{\in Z} : \mathcal{K} \models \langle A, \phi \rangle_{\mathcal{K}}$
- ... where A is an argument and $\mathcal{K} \models \langle A, \phi \rangle_{\mathcal{K}}$ is defined as in (13).¹⁹

6 General discussion and conclusion

The contributions of this paper are twofold. First, I've offered some evidence that the salient-argument felicity condition on *must* doesn't arise due to pragmatic competition with the bare assertion. This isn't to suggest that bare ϕ and *must* ϕ aren't alternatives. Rather, I argue that the argumentation problem won't be solved with a story that makes reference to counterfactual reasoning about alternative forms.

However, my solution is not entirely semantic in nature. While 'argumentation' is part of the semantic representation of *must* and *might*, context - in particular the contextual relevance of the speaker's reasoning - determines whether or not a particular argument must be identified by the listener for the modal to be felicitous.

¹⁹These entries suggest that *certain* and *possible* aren't semantic duals. \neg *Certain*($\neg\phi$) is true in Z iff there is no argument A that justifies $\neg\phi$ in all of Z 's argument systems. However, the argument systems in which A doesn't justify $\neg\phi$ could still be ones where some alternative argument A' does, so it's possible that \neg *certain*($\neg\phi$) is true even if there are no argument systems that support ϕ . So \neg *certain*($\neg\phi$) can be true even if *possible*(ϕ) is false.

A technical solution that preserves the duality assumption would be to simply have the existential quantifier of the argument occur inside the scope of the quantifier over argument systems:

- Certain ϕ is true in context Z iff $\forall \mathcal{K}_{\in Z} : \exists A : \mathcal{K} \models \langle A, \phi \rangle_{\mathcal{K}}$
- Possible ϕ is true in context Z iff $\exists \mathcal{K}_{\in Z} : \exists A : \mathcal{K} \models \langle A, \phi \rangle_{\mathcal{K}}$

Now, \neg *certain*($\neg\phi$) is true in Z when there's some argument system in Z in which no argument justifies $\neg\phi$. Assuming that Z is nonempty and each argument system contained in Z has an argument that justifies either ϕ or $\neg\phi$, this is equivalent to saying that there's an argument system where ϕ is justified - exactly the meaning we ascribe to *possible* ϕ . But strictly speaking, this new semantics allows for *certain* ϕ to be true if there are many potential arguments - none of them particularly likely in Z - for ϕ (as long as there are no potential arguments for $\neg\phi$). Lastly, these entries don't account for *certain* and *possible*'s gradability (e.g. *very certain* or *more possible (than ever)* are commonplace, as Klecha 2012, Klecha 2014, and Lassiter 2017 discuss).

A current weakness of the analysis is its reliance on a notion of ‘contextual relevance’ that should be more rigorously characterized. Crucially, I must assume that discourses can be structured not only by inquiry - e.g. as Questions Under Discussion (QUDs, Roberts 2012) - but also by the reasoning and evidence that helps to settle that inquiry. A more precise notion of context that dovetails with my semantics of epistemic modals is something I must leave to future work.

References

- Bürkner, P.-C. 2018. Advanced Bayesian multilevel modeling with the R package brms. *The R Journal* 10.395–411.
- Copley, B. 2006. What should ‘should’ mean? In *Proceedings of the Language Under Uncertainty Workshop, Kyoto University*.
- Davies, M. 2009. The 385+ million word corpus of contemporary american english (1990–2008+): Design, architecture, and linguistic insights. *International journal of corpus linguistics* 14.159–190.
- Degen, J., J. T. Kao, G. Scontras, & N. D. Goodman. 2015. A cost-and information-based account of epistemic must. *CUNY Sentence Processing Conference (Poster)*.
- Degen, J., A. Trotzke, G. Scontras, E. Wittenberg, & N. D. Goodman. 2019. Definitely, maybe: A new experimental paradigm for investigating the pragmatics of evidential devices across languages. *Journal of Pragmatics* 140.33–48.
- Del Pinal, G., & B. Waldon. 2019. Modals under epistemic tension. *Natural Language Semantics* 27.135–188.
- Karttunen, L. 1972. Possible and must. In *Syntax and Semantics 1*, ed. by J. Kimball.
- Katzir, R. 2007. Structurally-defined alternatives. *Linguistics and Philosophy* 30.669–690.
- Klecha, P. 2012. Positive and conditional semantics for gradable modals. In *Proceedings of Sinn und Bedeutung*, volume 16, 363–376.
- Klecha, P. 2014. *Bridging the divide: Scalarity and modality*. University of Chicago dissertation.
- Kratzer, A. 1991. Modality. In *Semantics: An international handbook of contemporary research*, ed. by A. von Stechow & D. Wunderlich, 639–650. De Gruyter Mouton.
- Lasnik, P. 2005. Context dependence, disagreement, and predicates of personal taste. *Linguistics and Philosophy* 28.643–686.
- Lassiter, D. 2016. Must, knowledge, and (in)directness. *Natural Language Semantics* 24.117–163.
- Lassiter, D. 2017. *Graded modality: Qualitative and quantitative perspectives*. Oxford University Press.
- Malamud, S. A., & T. Stephenson. 2014. Three Ways to Avoid Commitments: Declarative Force Modifiers in the Conversational Scoreboard. *Journal of Semantics* 32.275–311.
- Mandelkern, M. 2019. What ‘must’ adds. *Linguistics and Philosophy* 42.225–266.
- Ricciardi, G., R. Ryskin, & E. Gibson. 2020. Epistemic ‘must p’ is literally a strong statement. *CUNY Sentence Processing Conference (Poster)*.
- Roberts, C. 2012. Information structure: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics* 5.6–1.
- Simari, G. R., & R. P. Loui. 1992. A mathematical treatment of defeasible reasoning and its implementation. *Artificial Intelligence* 53.125–157.
- Stephenson, T. 2007. Judge dependence, epistemic modals, and predicates of personal taste. *Linguistics and Philosophy* 30.487–525.
- Stone, M. 1994. The reference argument of epistemic must. In *Proceedings of the International Workshop on Computational Semantics*.
- Swanson, E. 2006. *Interactions with context*. Massachusetts Institute of Technology dissertation.
- von Fintel, K., & A. S. Gillies. 2010. Must . . . stay . . . strong! *Natural Language Semantics* 18.351–383.
- von Fintel, K., & A. S. Gillies. in press. Still going strong. *Natural Language Semantics*.
- Waldon, B. in press. Linguistic interpretation as inference under argument system uncertainty: the case of epistemic must. In *Proceedings of Probability and Meaning (PaM2020)*.
- Yalcin, S. 2016. Modalities of normality. In *Deontic Modality*, ed. by N. Charlow & M. Chrisman. Oxford University Press.