

Linguistic interpretation as inference under argument system uncertainty: the case of epistemic *must*

Brandon Waldon

Stanford University Department of Linguistics, 450 Jane Stanford Way

Stanford, CA 94305 USA

bwaldon@stanford.edu

Abstract

Modern semantic analyses of epistemic language (incl. the modal *must*) can be characterized by the ‘credence assumption’: speakers have full certainty regarding the propositions that structure their epistemic states. Intuitively, however: a) speakers have graded, rather than categorical, commitment to these propositions, which are often never fully and explicitly articulated; b) listeners have higher-order uncertainty about this speaker uncertainty; c) *must* ϕ is used to communicate speaker commitment to some conclusion ϕ and to indicate speaker commitment to the premises that condition the conclusion. I explore the consequences of relaxing the credence assumption by extending the argument system semantic framework first proposed by Stone (1994) to a Bayesian probabilistic framework of modeling pragmatic interpretation (Goodman and Frank, 2016).¹

1 Introduction

Natural language contains a variety of means for expressing one’s epistemic state. The best-studied of these in the semantics literature are the epistemic modal auxiliaries *must* and *may/might*:

- (1) a. Ann: *Where is Peter?*
- b. Mary: *He {may/might/must} be in his office.*

There is broad agreement in the literature that Mary’s response in (1b) is comprised in part by a conclusion - *Peter is in his office* - the ‘prejacent’ over which the modal takes semantic scope. Additionally, the consensus is that the epistemic modal expresses a connection between the prejacent and a set of salient premises - most commonly, things that are known and/or assumed by the speaker and/or her interlocutor. Roberts (2019)

notes that the details beyond these points of agreement are matters of debate; in particular, theoreticians disagree over the following two questions:

1. How do we specify the premises - the body of information, assumptions, or other contextually-supplied propositions which condition a modalized statement?
2. In what way are the premises related to the conclusion ϕ encoded as the prejacent of a modalized statement?

A well-discussed desideratum of a successful theory of epistemic modality is that it should provide an understanding of the perceived weakness of *must*. An observation going back to Karttunen (1972) is that modalized statements of the form *must* ϕ appear to mark weak speaker commitment to the prejacent compared to the unmodalized counterpart, ‘bare’ ϕ . The observation stems from consideration of contexts such as (2):

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

Answers to Roberts’ original questions fall into three categories. Restricted quantificational accounts (Kratzer, 1991) posit that the set of premises includes propositions known to a contextually-salient individual (most often the speaker) or group of individuals (containing the speaker and her interlocutor), as well as a contextually-specified set of assumptions which further restrict the space of epistemic possibilities. *Must/might* ϕ quantify universally and existentially over this space, respectively: *must* expresses that the conclusion is true at all possible worlds where the known and assumed propositions are true; *might* expresses that the conclu-

¹I gratefully acknowledge Cleo Condoravdi, Judith De- gen, Atticus Geiger, Daniel Lassiter, Christopher Potts, three PaM reviewers, and Stanford’s Construction of Meaning workshop for valuable feedback. All errors are mine.

sion is true at at least one of those worlds. Perceived weakness of *must* is accounted for on this analysis because *must* ϕ - unlike bare ϕ - allows for the possibility that $\neg\phi$ is true in worlds compatible with the known propositions (but incompatible with the assumed ones). In contrast, unrestricted quantificational accounts (von Fintel and Gillies, 2010) posit that *must* quantifies over a space of epistemic possibilities that is unconstrained by contextually-salient assumptions. On this approach, *must* ϕ is incompatible with $\neg\phi$, and *must*'s infelicity in (2) is a consequence of a violation of independently-stipulated felicity conditions.² Finally, probabilistic accounts (Swanson, 2006; Lassiter, 2016) vary in their commitments regarding how to specify the premises but consider *must/might* to be operators which take as their input the premises and output a statement about the likelihood of the truth of the prejacent.

All of the above approaches can be characterized by what I call the credence assumption - that is, that however we specify the premises and their relation to the conclusion, speakers have full certainty about the premises upon which they can rely for the purposes of inference in a given context. This assumption is desirable from the standpoint of analytic simplicity; moreover, it provides a way of analyzing modal disagreement, as in (3):

- (3) a. Ann: *It might be raining outside.*
 b. Mary: *No, it cannot be raining outside!*

On the credence assumption, Ann makes a statement regarding the possibility of rain given a set of known and/or assumed propositions. Mary assesses this statement and disagrees: she has a different (yet also deterministic) understanding of the premises operable in this discourse.³

Intuitively, however, speakers' epistemic states are much more complex than the credence assumption allows. Namely, these states involve graded, rather than deterministic, commitments to propositions that are often never fully and explicitly articulated by speakers who produce statements of the form *must/might* ϕ . Listeners in turn have uncertainty about the premises which their interlocutors deem to be relevant for the purposes of inference and deliberation; indeed, *must/might* ϕ is informative not just because it

²von Fintel and Gillies (2010), for example, contend that *must* ϕ presupposes that ϕ has not yet been settled in context.

³Or perhaps Mary agrees with Ann on the premises but disagrees regarding their relationship to the conclusion.

conveys speaker's epistemic commitment to the prejacent but also because it conveys something about the speaker's underlying knowledge and assumptions about the world, and about how she is likely to use available information in the future.

I develop a quantitative framework of modeling interpretation of *must* that relaxes the credence assumption. In doing so, I offer a formal means of representing how these constructions can be informative with respect to speaker commitment to the conclusion as well as with respect to the premises that the speaker believes are operable in context. My point of departure is the argument system semantic framework of Stone (1994), followed by a probabilistic enrichment of that framework rooted in a Bayesian understanding of linguistic inference (Goodman and Frank, 2016). On this approach, communication proceeds between agents who are uncertain about what premises can (and should) be relied upon for the purposes of present and future inference and deliberation. Interlocutors align this uncertainty in part via communicative exchange. Finally, I consider implications of this approach for our understanding of conversational dynamics and the common ground.

2 Argument system semantics

Stone (1994) observes that in contexts such as (1), the *must*-variant of Mary's response is infelicitous if the context does not make clear (to Ann) the basis on which Mary's conclusion is made: Mary's conclusion about Peter may come from the fact that Mary has ruled out all possible other places Peter could be, or perhaps it is 3pm on a Tuesday and Peter is always in his office at that time. If Ann cannot recover Mary's argument in support of the conclusion, then *must* is infelicitous in (1). On Stone (1994)'s semantics, *must* ϕ is true iff a (possibly defeasible) argument A - made somehow salient in the context - justifies concluding ϕ given an argument system \mathcal{K} . I recapitulate the relevant details of Stone's analysis below.⁴

2.1 Formal preliminaries

Let \mathcal{K} be an argument system, comprised of a set of established propositions K (ground formulae

⁴I focus on Stone's argument system semantics because his formalism provides a way of verifying the relationship between a conclusion and the premises that condition it. This is crucial for my analysis, which captures how listeners infer speaker beliefs about premises having only observed conclusions asserted by the speaker. But similar results could be achievable with other semantic 'backends'.

K_C and logical rules of inference K_N) and a set of defeasible inferential rules Δ . Arguments for ground formulae are defined as follows:⁵

Definition 1: A set T of instantiations of elements of Δ is an ARGUMENT for h ($\langle T, h \rangle_{\mathcal{K}}$) iff: (1) $K \cup T \vdash h$; (2) $K \cup T \not\vdash \perp$; and (3) for no $T' \subset T$, $K \cup T' \vdash h$.

The first clause of Definition 1 specifies that an argument for a ground formula (comprised of elements of Δ), added to K , entails the formula; the second specifies that the argument must be consistent with K ; the third specifies that the argument must be minimal. Stone also introduces the notion of a sub-argument: an argument which can be computed from the premises of another argument:

Definition 2: $\langle S, j \rangle_{\mathcal{K}}$ is a SUBARGUMENT of $\langle T, h \rangle_{\mathcal{K}}$ if and only if $S \subseteq T$.

Stone emphasizes that subarguments of $\langle T, h \rangle_{\mathcal{K}}$ need not play a role in concluding h . Rather, the set of subarguments for h include all arguments which can be generated from T given argument system \mathcal{K} . This means that counterarguments to $\langle T, h \rangle_{\mathcal{K}}$ can do so by targeting not only subarguments necessary to conclude h from T but any inference generated from T given \mathcal{K} .

Definition 3: $\langle T_1, h_1 \rangle_{\mathcal{K}}$ COUNTERARGUES $\langle T_2, h_2 \rangle_{\mathcal{K}}$ at $\langle T, h \rangle_{\mathcal{K}}$ if and only if $\langle T, h \rangle_{\mathcal{K}}$ is a subargument of $\langle T_2, h_2 \rangle_{\mathcal{K}}$ and $K \cup \{h, h_1\} \vdash \perp$.

A counterargument defeats an argument if the former is more specific - if it “takes more particulars of the context into consideration” (1994: 6).

Definition 4: $\langle T_1, h_1 \rangle_{\mathcal{K}}$ is more SPECIFIC than $\langle T_2, h_2 \rangle_{\mathcal{K}}$ if and only if: (1) for all ground formulae e , if $K_N \cup \{e\} \cup T_1 \vdash h_1$ but $K_N \cup \{e\} \not\vdash h_1$, then $K_N \cup \{e\} \cup T_2 \vdash h_2$; and (2) there is some ground e such that $K_N \cup \{e\} \cup T_2 \vdash h_2$, $K_N \cup \{e\} \cup T_1 \not\vdash h_1$, and $K_N \cup \{e\} \not\vdash h_2$.

Definition 5: $\langle T_1, h_1 \rangle_{\mathcal{K}}$ DEFEATS $\langle T_2, h_2 \rangle_{\mathcal{K}}$ if $\langle T_1, h_1 \rangle_{\mathcal{K}}$ counterargues $\langle T_2, h_2 \rangle_{\mathcal{K}}$ at $\langle T, h \rangle_{\mathcal{K}}$ and $\langle T_1, h_1 \rangle_{\mathcal{K}}$ is more specific than $\langle T, h \rangle_{\mathcal{K}}$

The first clause of Definition 4 states that the conclusion of the less specific argument h_2 must be entailed by the argument system coupled with the argument’s defeasible premises T_2 , provided the argument system is one in which the more specific argument’s conclusion h_1 only follows with the addition of its premises T_1 . The second clause states that there must be some argument system

which entails the conclusion of the less specific argument h_2 on the basis of the premises T_2 but is inconsistent with the conclusion of the more specific argument h_1 on the basis of premises T_1 .

An argument system justifies an argument “whenever [the argument] has no counterarguments which are not themselves defeated” (1994: 6). To formalize this, Stone introduces the concepts of supporting arguments and interfering arguments, defined inductively to capture the fact that for an argument to be justified it must not be defeated at any level of sub-argumentation.

Definition 6: All arguments are level 0 supporting and interfering arguments.

- An argument $\langle T_1, h_1 \rangle_{\mathcal{K}}$ is a level $(n + 1)$ supporting argument if and only if no level n interfering argument counters it at any of its subarguments.
- An argument $\langle T_1, h_1 \rangle_{\mathcal{K}}$ is a level $(n + 1)$ interfering argument if there is no level n interfering argument which defeats it.

Definition 7: An argument $\langle T, h \rangle_{\mathcal{K}}$ JUSTIFIES h in \mathcal{K} if and only if there is some m such that for all $n \geq m$, $\langle T, h \rangle_{\mathcal{K}}$ is a level n supporting argument. h is justified in \mathcal{K} if some $\langle T, h \rangle_{\mathcal{K}}$ justifies it in \mathcal{K} .

2.2 Illustration

First, let \mathcal{K}_0 be an argument system consisting of ground formulae K_{0C} , logical rules K_{0N} , and defeasible rules Δ_0 .⁶ Assume K_{0N} contains a forward chain inferential rule, and Δ_0 consists of the following defeasible rules about matches and heat:

- A: $\text{match}(x) \wedge \text{strike}(x) > \text{lit}(x)$
 B: $\text{match}(x) \wedge \text{strike}(x) \wedge \text{wet}(x) > \neg \text{lit}(x)$
 C: $\text{lit}(x) > \text{hot}(x)$

Let K_{0C} contain two ground formulae $\text{match}(m1)$ and $\text{strike}(m1)$. By Definition 1, we generate two arguments, A_1 and A_2 :

- $A_1: \langle \{A\}, \text{lit}(m1) \rangle_{\mathcal{K}_0}$
 $A_2: \langle \{A, C\}, \text{hot}(m1) \rangle_{\mathcal{K}_0}$

Now, consider Stone’s semantics for *must*:

- (4) *Must* ϕ is true in \mathcal{K} iff $\mathcal{K} \models \langle A, \phi \rangle_{\mathcal{K}}$

That is, *Must* ϕ is true if a contextually-salient argument A justifies concluding ϕ in a given argument system. Given \mathcal{K}_0 , (5) is true:

- (5) The match must have lit.

⁵All definitions below can be found in Stone (1994): p. 6.

⁶This example is based largely on one from Stone (1994).

Note that the semantics of *must* requires (in the case of 5) that the argument (A_1) is contextually-salient; otherwise, *must* is undefined. Assuming this condition is met, we can verify the truth of (5) by considering, by Definition 7, whether concluding $\text{lit}(m1)$ from A_1 is justified in \mathcal{K}_0 . It is: the only arguments that could interfere would have as their conclusion $\neg\text{lit}(m1)$, but these arguments cannot be generated from \mathcal{K}_0 because $\text{wet}(x)$ is not in \mathcal{K}_{C_0} (and Rule B cannot be invoked).

Now consider a second argument system \mathcal{K}_1 , which differs minimally from \mathcal{K}_0 in that $\text{wet}(m1)$ is an additional ground formula. Thus, A_3 is generated in addition to A_1 and A_2 :

$A_3: \langle \{B\}, \neg\text{lit}(m1) \rangle_{\mathcal{K}_1}$

In \mathcal{K}_1 , (5) is false. Note first that the only argument from which $\text{lit}(m1)$ can be concluded is A_1 . Thus, as in \mathcal{K}_0 , (5) can only be true if A_1 is justified. It is not in \mathcal{K}_1 : A_3 defeats A_1 because the former is more specific than the latter.

2.3 Interim discussion

Stone’s system provides a straightforward account of *must*’s perceived weakness, if we can assume that direct observation of h adds h to the set of ground formulae by default. Consider Definition 1: an argument for h in \mathcal{K} must have as its premises the minimal set of defeasible rules of inference which - coupled with the set of established ground formulae and the logical rules of inference in \mathcal{K} - entails h . If h is already in the set of ground formulae, then the minimal set of required premises is empty. The prediction is that there is no argument A that can meet the definedness conditions of *must* ϕ if ϕ is already established in \mathcal{K} .

Argument systems and speaker epistemic states are assumed to be one and the same on this analysis, meaning that this analysis can be characterized by the credence assumption. We might imagine one speaker whose epistemic state is akin to argument system \mathcal{K}_0 , and another whose state is akin to \mathcal{K}_1 . On this analysis, it is clear why these two agents might disagree over (5): the statement is true given the former argument system and false given the latter. Below, I explore the properties of an extension that relaxes the credence assumption.

3 Probabilistic argument system semantics

In the context of Stone’s analysis, relaxing the credence assumption amounts to revising our as-

sumptions regarding the speaker’s relationship to \mathcal{K} . I define a space of possible argument systems Z , which I allow to vary according to their ground formulae and defeasible rules of inference. I assume that speakers are uncertain as to what precise argument system is the relevant one for the purposes of inference and decision making in context. That is, there may be some uncertainty as to whether particular ground formulae can be taken to be true at the world of evaluation, or there may be uncertainty as to whether certain defeasible rules of inference may (or should) be employed in a particular context. I assume that Z is specified such that the ground formulae that the speaker considers likely to be true are in many (but not perhaps not all) of the elements of Z ; the same is assumed modulo the defeasible rules of inference.⁷

We can then define speaker commitment to a proposition ϕ on the basis of some argument A as the likelihood that $\langle A, \phi \rangle$ justifies ϕ given possible argument systems in Z . *Must* ϕ , then, is a comment on this likelihood value: if the likelihood exceeds a certain contextually-specified threshold, then the statement is true:

$$(6) \quad \text{Must } \phi \text{ is true in } Z \text{ iff } P(\langle A, \phi \rangle)_Z > \theta, \\ \text{where } P(\langle A, \phi \rangle)_Z = \frac{\sum_{\mathcal{K} \in Z} \mathcal{K}_{F(\langle A, \phi \rangle)} \mathcal{K}}{|Z|}$$

Speakers produce *must* ϕ to convey their degree of belief that ϕ is a valid conclusion on the basis of an argument A , given their argument system uncertainty. But importantly, the precise nature of this argument system uncertainty - the precise value of Z - is not transparent to the listener: the listener has prior beliefs about possible values of Z that are updated according to the conclusions that a speaker draws (and argumentation that she employs to draw those conclusions) in a particular context. Observation of *must* ϕ , then, allows the listener to update her uncertainty about the speaker’s Z distribution.

The speaker’s production of *must* ϕ is determined by a utility function of utterances given intended meanings that balances informativity against production cost. Following Goodman and Frank (2016), the model of a pragmatic speaker S_1 is defined partly in reference to a literal L_0 listener whose interpretations are a function of utterances’ literal truth/falsity given possible intended

⁷That is, on this analysis, every individual argument system is assumed to have uniform probability. Alternatively, as a reviewer suggests, one could suppose that some elements of Z are more likely than others a priori (and that the truth conditions of *must* ϕ are sensitive to this non-uniform prior).

meanings. The space of possible meanings that the speaker could try to convey are possible valuations of the speaker’s Z distribution (from which - by 6 - the speaker’s commitment to ϕ can be computed).

The literal L_0 listener, then, can be modeled as a conditional probability distribution over possible valuations of Z given observation of some utterance u and a contextually-supplied value for the probability threshold θ . Following Lassiter and Goodman (2013), who model interpretation of gradable adjectives using a threshold-based semantics, I assume that this θ variable is passed from L_0 and eventually estimated by the pragmatic L_1 listener (defined below) from a prior distribution over values of θ .

$$P_{L_0}(Z|u, \theta) = P_{L_0}(Z|\llbracket u \rrbracket^\theta = 1) \times P(Z)$$

The pragmatic speaker selects utterances to convey intended meanings according to their contextual informativeness for L_0 as well as the cost of utterance production. Below, α is a speaker optimality parameter, and C is a cost function defined for all possible utterance choices: all else equal, the greater $C(u)$, the lower the probability that S_1 selects u to convey a particular message.

$$P_{S_1}(u|Z, \theta) \propto \exp(\alpha \times \log(P_{L_0}(Z|u, \theta)) - C(u))$$

The pragmatic listener L_1 ’s interpretations of utterances are a function of expected behavior of S_1 , as well as prior expectations about the likelihood of different possible meanings and prior expectations about the threshold value θ :

$$P_{L_1}(Z, \theta|u) \propto P_{S_1}(u|Z, \theta) \times P(Z, \theta)$$

Thus, interpretation of *must* is a joint inference about the state of the world (speaker beliefs regarding the justifiability of concluding ϕ) and the value of a semantic threshold variable θ . These speaker beliefs - $P(\langle A, \phi \rangle_Z)$ - can be computed given values the argument variable supplied categorically by the context (A) and one additional variable (Z) which is inferred under uncertainty.

3.1 Illustration

For this illustration, I assume that the listener has uniform prior beliefs over the threshold value θ and that she considers two possible utterance production choices: *must* - whose truth conditions are as in (6) - and a trivially true null message.⁸

⁸I follow Lassiter and Goodman (2013) in introducing this null utterance choice, which is an implementational necessity in the absence of utterance alternatives. Adding plausible linguistic alternatives to the model - including *might* ϕ and bare ϕ - does not drastically alter the patterns presented here.

The listener assumes that the speaker has full certainty about the following features of the argument system: the ground formulae (consisting of propositions `match(m1)`, `strike(m1)`, and `wet(m1)`), the logical rules of inference (including a forward chain operation), and a subset of the defeasible rules of inference (i.e. the listener assumes that Rule A features in every candidate argument system considered by the speaker). However, there are two other inferential rules - Rules B, and C from above - the status of which the speaker is uncertain: elements of Z may individually feature one, both, or neither of these rules.

For this illustration, assume that the listener has observed the speaker utter (5). Intuitively, this utterance conveys a high degree of speaker commitment to the prejacent (`lit(m1)`), but it should also convey something to the listener about the speaker’s argument system uncertainty: since it is established that the speaker recognizes that `wet(m1)`, in uttering (5) the speaker has signalled that she finds it unlikely that Rule B is a relevant premise in this context.

The pragmatic listener must infer the value of Z under uncertainty; that is, she will not know the precise proportion of elements of Z that contain inferential Rules B and/or C (or neither). In other words, let β be the speaker’s degree of belief that Rule B is in the contextually-relevant argument system; and let γ stand in for speaker beliefs about Rule C. The pragmatic listener updates her beliefs about the values of β and γ by observing the speaker’s utterance production choices in context, in addition to inferring the value of the threshold θ . For this illustration, I assume uniform prior beliefs over values for β , γ , and θ and make the simplifying assumption that $C(\textit{must})$ is equal to 1 while the null message has zero cost.⁹ I arbitrarily set the optimality parameter α to 4.

In the computational implementation of this example, 10,000 samples are drawn from $P_{L_1}(\beta, \gamma, \theta|\textit{must}(\textit{lit}(\textit{m1})))$ using Markov Chain Monte Carlo sampling, with the assumption that the contextually-salient argument A is A_1 .¹⁰ Marginal posterior distributions over values for the inferred parameters are presented in Figure

⁹The prior distributions over values for β , γ , and θ are discrete distributions with uniform probability mass on 11 evenly-spaced values on the interval $[0, 1]$.

¹⁰The implementation was programmed using WebPPL (Goodman and Stuhmüller, 2014). Code is available at <https://github.com/bwaldon/probmust>.




Figure 1: Marginal posterior distributions over values of β , listener beliefs about the speaker’s expectation that Rule B is in the contextually-relevant argument system; γ , listener beliefs about the speaker’s expectation that Rule C is in this state; θ , the threshold for *must*, and speaker commitment to *lit(m1)* on the basis of inferential rule A, calculated by approximating a posterior distribution over values of Z from posterior values of β and γ . Degree of speaker commitment is anti-correlated with β and not correlated with γ .

1. As a sanity check, we see that the posterior over values of γ is effectively uniform. This is exactly what is to be expected, as the inclusion of Rule C in the argument system has no bearing on the justifiability of concluding *lit(m1)*; thus, the speaker’s production of (5) is not informative for the listener regarding the status of Rule C. However, the posterior over values of β suggests that the listener has learned something regarding the speaker’s beliefs about Rule B.

Recall that the presence of this inferential rule in the argument system has the consequence that $\langle [match(x) \wedge strike(x) > lit(x)], lit(m1) \rangle$ is not justified (given our assumptions about the ground formulae and possible rules of inference from above). But *must-lit(m1)* was asserted on the basis of argument $match(x) \wedge strike(x) > lit(x)$; thus, after observing the speaker produce (5), the listener considers it relatively unlikely that the speaker expects Rule B - $match(x) \wedge strike(x) \wedge wet(x) > \neg lit(x)$ - to be a relevant inferential rule.

4 Discussion and conclusion

On this picture, disagreement functions differently than on analyses characterized by the credence assumption. On that assumption, we could understand disagreements over *must* ϕ as stemming from interlocutors’ differences regarding their (deterministic) beliefs about the status of the premises or regarding the relationship of the premises to ϕ . The probabilistic enrichment explored here makes the story slightly more complicated. Consider the illustration above: a listener who hears a speaker utter (5) in this context is likely to disagree with that speaker, if the listener’s own uncertainty in-

volves high expectation that Rule B is relevant for the purposes of inference (and hence the listener has relatively low commitment to the pre-jacent, *the match is lit*). But what is the source of the disagreement? It cannot be that the listener knows definitively that she and the speaker have drastically different expectations regarding what inferential premises can be relied on in this context. However, the speaker’s production of (5) is highly suggestive of such a difference: it is quite likely that the speaker puts relatively little weight in the chance that wet matches will light, even when struck. As a consequence, it is quite likely that the speaker has a high degree of belief that the match is lit. Disagreement, then, is triggered by the listener being fairly certain that her argument system uncertainty - her internal Z distribution - is substantially different from her interlocutor’s.¹¹

This suggests a way of understanding context and communicative exchange that complements the conventional “common ground” approach of Stalnaker (2002), whereby context records the propositions that interlocutors accept (‘treat as true’), and communicative exchange involves proposals to update this common ground via addition of new propositions. In particular, my analysis suggests a means of formally modeling another layer of the context concerned with the uncertainty that interlocutors bring to bear on propositions not necessarily treated as categorically true. Epistemic linguistic constructions (e.g. *must*) facilitate coordination of this uncertainty between interlocutors, by communicating properties of this uncertainty from a particular epistemic vantage point.

¹¹A more precise understanding of modal disagreement in this framework - for example, how do we quantify the conditions giving rise to disagreement? - is left to future work.

References

- Kai von Fintel and Anthony S. Gillies. 2010. *Must . . . stay . . . strong!* *Natural Language Semantics*, 18(4):351–383.
- Noah D Goodman and Michael C Frank. 2016. Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences*, 20(11):818–829.
- Noah D Goodman and Andreas Stuhlmüller. 2014. The Design and Implementation of Probabilistic Programming Languages. <http://dippl.org>. Accessed: 2020-6-17.
- Lauri Karttunen. 1972. Possible and must. In J. Kimball, editor, *Syntax and Semantics*, volume 1, pages 1–20. Academic Press, New York.
- Angelika Kratzer. 1991. Modality. In Arnim von Stechow and Dieter Wunderlich, editors, *Semantics: An International Handbook of Contemporary Research*, volume 7, pages 639–650. de Gruyter, Berlin.
- Daniel Lassiter. 2016. *Must, knowledge, and (in)directness*. *Natural Language Semantics*, 24(2):117–163.
- Daniel Lassiter and Noah D Goodman. 2013. Context, scale structure, and statistics in the interpretation of positive-form adjectives. In *Proceedings of Semantics and Linguistic Theory*, volume 23, pages 587–610.
- Craige Roberts. 2019. The character of epistemic modality: Evidential indexicals. *Ms., Ohio State University*.
- Robert Stalnaker. 2002. Common ground. *Linguistics and Philosophy*, 25(5/6):701–721.
- Matthew Stone. 1994. The reference argument of epistemic must. In *Proceedings of the International Workshop on Computational Semantics*, volume 1, pages 181–190.
- Eric Swanson. 2006. *Interactions with context*. Ph.D. thesis, Massachusetts Institute of Technology.