

MIT Open Access Articles

Must . . . stay . . . strong!

The MIT Faculty has made this article openly available. ***Please share*** how this access benefits you. Your story matters.

Citation: von Fintel, Kai, and Anthony S. Gillies. "Must . . . Stay . . . Strong!" *Natural Language Semantics* 18.4 (2010) : 351-383.

As Published: <http://dx.doi.org/10.1007/s11050-010-9058-2>

Publisher: Springer Netherlands

Persistent URL: <http://hdl.handle.net/1721.1/65909>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike 3.0



Must ... Stay ... Strong!*

Kai von Fintel
MIT

Anthony S. Gillies
Rutgers University

This is a postprint of a paper that appeared in *Natural Language Semantics*:

von Fintel, Kai & Anthony S. Gillies. 2010. *Must ... stay ... strong!* *Natural Language Semantics* 18(4). 351–383. doi:10.1007/s11050-010-9058-2.

1 The Observation, the Mantra, and the Plan

Some epistemic modal expressions are stronger than others. By asserting any expression in the list (1), a speaker stakes herself to a stronger claim than she would by asserting the next in the list:

- (1) a. It must be raining.
- b. It is almost certainly raining.
- c. It, presumably, is raining.
- d. It might be raining.

That much is clear.¹ But what about the relation between one of these modalized claims and the bare prejacent *it is raining*? Where does that belong on the list? Is it stronger or weaker than “strong” epistemic necessity modals like *must* in English? That is less clear and is our focus in this paper.

The basic quantificational treatment of modality makes a clear prediction. Since *must* is English’s way of representing an operator of epistemic necessity, it quantifies over possibilities compatible with what is known, saying that all of them are possibilities in which the prejacent is true. Whatever is known is true, and so the actual situation is always among those compatible with what is known. Thus, for any ϕ , *must* ϕ entails ϕ . The strong modalized claim is stronger than the bare prejacent. Thus, ϕ ’s proper place on the list is below

* This is the fourth installment in our trilogy of papers on epistemic modality. We would like to thank audiences at NYU, MIT, the University of Michigan, the University of Düsseldorf, the University of Texas at Austin, the Second Arché Contextualism and Relativism Workshop at St. Andrews, the University of British Columbia, and Rutgers University for discussion and comments, and William Shatner for inspiring the title. Special thanks to Chris Potts, Craig Roberts, and two anonymous reviewers for their very useful comments; thanks as well to Yimei Xiang for comments and questions.

¹ The importance of the fact that modal force is a graded notion was of course emphasized by Kratzer in seminal work (1981; 1991).

must ϕ 's.²

But this prediction is hard to square with a very basic intuition. When one considers which of the answers to the question in (2) conveys more confidence and inspires more confidence, it is natural to feel that the strongly modalized claim is less forceful:

- (2) Where are the keys?
 - a. They are in the kitchen drawer.
 - b. They must be in the kitchen drawer.

Since Karttunen (1972) was the first semanticist who pointed out the issue, we will call this “Karttunen’s Problem”: modal semantics predicts that (2b) is a stronger answer to the question than (2a), but naïve intuition goes the other way.

Confronted with Karttunen’s Problem, semanticists have reacted with an overwhelming consensus that the meaning of epistemic *must* needs to be weaker than classically predicted and weaker than that of the bare preja-cent — a consensus that has Mantra status. Here are four early-ish quotes:

Karttunen (1972) writes that “Intuitively, (3b) makes a weaker claim than (3a)”:

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

Groenendijk & Stokhof (1975: 69) write: “A statement like (4a) is weaker than (4b). (4b) expresses more conviction on the part of the speaker than (4a) does”:

- (4) a. John must be at home.
 - b. John is at home.

Lyons (1977: 808) puts the point less succinctly:

Although it might appear that a statement is strengthened by putting the proposition that it expresses within the scope of the operator of epistemic necessity, this is not so, as far

² This prediction is so basic to standard modal semantics that we will relegate a slightly more detailed rehearsal of how it comes about to Appendix A.

as the everyday use of language is concerned. It would be generally agreed that the speaker is more strongly committed to the factuality of *It be raining* by saying *It is raining* than he is by saying *It must be raining*.

Kratzer (1991) says “It has often been observed that I make a stronger claim in uttering (5a) than in uttering (5b)”:

- (5) a. She climbed Mount Toby.
- b. She must have climbed Mount Toby.

We could multiply such citations at will.

Usually, the weakness intuition is simply announced without much argument, just as we did ourselves above when talking about (2). There is, though, another well-known fact about epistemic *must* that can be shown quite clearly and that is intertwined with the weakness claim in much of the literature: *must* carries an *evidential* signal, in particular it signals that the speaker has reached her conclusion via an indirect inference.

Billy is a weather enthusiast. Looking out the window seeing pouring rain, she can report by asserting (6a) but not (6b).

- (6) [Seeing the pouring rain]
- a. It’s raining.
- b. ??It must be raining.

If, instead, she sees people coming in from outside with wet umbrellas, slickers, and galoshes, then — even if she knows that rain is the only explanation — she can report with either the modalized claim or its bare prejacent. Either will do:

- (7) [Seeing wet rain gear and knowing rain is the only possible cause]
- a. It’s raining.
- b. It must be raining.

This is not the place for us to survey the field of evidentiality studies in any amount of detail (some relevant references are Aikhenvald 2004; Davis, Potts & Speas 2007; Faller 2002; de Haan 2001; Matthewson, Rullmann & Davis 2007; McCready & Ogata 2007; Rooryck 2001a,b; Willett 1988). Suffice it to say that evidential markers are expressions found in many languages that signal the source of evidence a speaker has for the prejacent claim. Evidentials

often come in a system of related meanings. A nice summary of the typical distinctions made in rich evidential systems across languages is given in Figure 1 from Willett (1988).

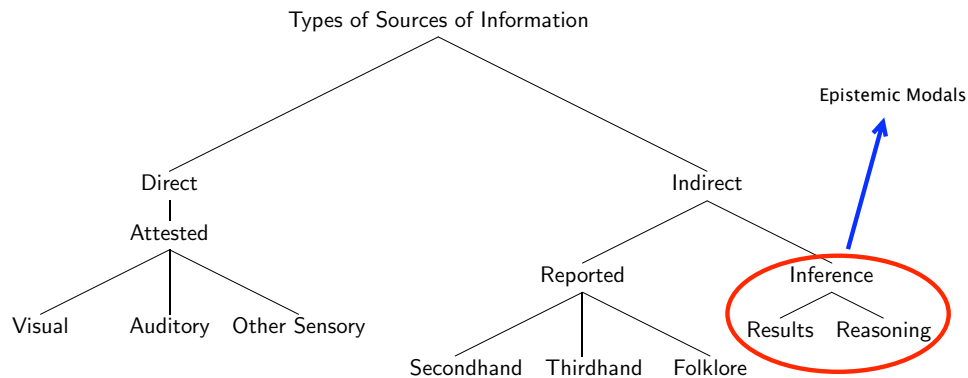


Figure 1 Willet's Taxonomy of Evidential Categories

The data in (6)/(7) suggest that epistemic modals are also evidential markers: they signal that the prejacent was reached through an inference rather than on the basis of direct observation or trustworthy reports. It appears that seen as evidentials, epistemic modals are markers of INDIRECT INFERENCE, that is the rightmost branch of Willet's system. Thus, we have enriched the graphic with an indication of where the evidential signal carried by epistemic modals fits in.^{3,4}

³ Note that epistemic modals do not cover the notion of indirect evidence derived from reports (the sister of indirect inference in Willet's system). Frank Jackson (pc) gave us a relevant scenario. When one reads in a book that the Battle of Hastings was fought in 1066, one's evidence is indirect; one does not observe the battle or anything like that. But it would be wrong to say "The Battle of Hastings must have been fought in 1066". Or, again, after reading the newspaper report about last night's Knicks game, one doesn't say "The Knicks must have lost again." The English pseudo-evidential *apparently* is more appropriate here.

⁴ It should be noted that the literature on evidentials often makes a strict distinction between epistemic modality and evidentiality, but the facts we discuss here indicate that this is too simplistic a position. See also Blain & Déchaine (2007) and Matthewson et al. (2007) for cross-linguistic evidence for the close connection between epistemic modality and evidentiality. This is not to say that there aren't evidential constructions that do not behave like standard epistemic modals; the evidential literature is full of such cases. It also doesn't mean that there couldn't be epistemic modals that are not evidentials, although we haven't found

So, the situation is this: we have strong reasons to think *must* expresses and inspires less confidence than it should if it had a standard necessity semantics; we have possibly even stronger reasons to think that uses of it signal that the speaker is basing their claim on indirect inferences. There is a wide variety of possible analyses to capture these facts. Mantra-based analyses propose to weaken the meaning of epistemic *must*. As we'll see they come in at least two variants, depending on whether they locate the weakness in the semantics or in the pragmatics of assertion. Mantra-based analyses also generally make the indirect evidential signal part of that weak treatment. What we will argue for, instead, is a combination of the evidential component with a truth-conditional semantics that keeps *must* as strong as it ever was.

Here, then, is the plan for this paper. We will survey two kinds of approaches that make *must*-statements weaker than the bare prejacent. One approach (historically much prior and recently revived) locates the weakness in the force of assertion rather than in the truth-conditional level. We will quickly dismiss this approach because it fails to get embedding facts right. The other approach (probably the standard way to go) operates on the truth-conditional level. We then argue that in fact, the semantics of *must* should remain strong. We do this by refuting the tenability of a weak semantics and then by working out a way of integrating the evidential component in a strong semantics.

2 Kant and Frege

The oldest approach to Karttunen's Problem was inaugurated quite a while ago. The claim is that *must* operates entirely outside the content dimension.⁵ At the level of propositional content, the story goes, the bare prejacent and the *must* statement are completely identical, and all that *must* contributes is a "comment" about the speaker's evidential basis for the prejacent proposition. This is the kind of analysis that is quite widespread in the descriptive linguistic literature on (epistemic) modality.⁶ The prevalence of this con-

any — a fact that is relevant to the question whether the evidential signal should be seen as part of the conventional meaning of epistemic modals, or whether it could be derived as a conversational implicature, see Section 5.

⁵ Parts of this section are based on material from our "Opinionated Guide" (von Fintel & Gillies 2007b). Portner (2009: esp. Section 4.2) ably reports and elaborates on these arguments and develops his own perspective. We highly recommend his book for further reading.

⁶ See, for example, the passages quoted by Papafragou (2006: 1688-1689).

ception can perhaps be traced back to the influence of Immanuel Kant, who wrote in his *Critique of Pure Reason* that “the modality of judgments is a very special function thereof, which has the distinguishing feature that it does not contribute to the content of the judgment” (1781: 74). This idea seems to have influenced both practicing linguists and a subset of logicians, including Gottlob Frege, who wrote in his *Begriffsschrift* that “[w]hat distinguishes the apodeictic from the assertoric judgment is that it indicates the existence of general judgments from which the proposition may be inferred — an indication that is absent in the assertoric judgment. By saying that a proposition is necessary I give a hint about the grounds for my judgment. But, since this does not affect the conceptual content of the judgment, the form of the apodeictic judgment has no significance for us” (1879: 5).

A frequently cited, more recent proponent of an analysis that locates an evidential signal in the comment dimension of *must* statements is Westmoreland (1995, 1998).⁷ According to Westmoreland, epistemic *must* “contributes the information that the propositional content of the sentence is inferred rather than known” (Westmoreland 1995: 699).⁸ Westmoreland goes further: while the bare prejacent is normally asserted in a strong sense (one could assume a knowledge norm of assertion, perhaps), epistemic *must* statements do not assert the underlying proposition: “an evidential sentence does not assert a proposition A sentence with a ‘hearsay’ marker presents a proposition as being second-hand information; a sentence with *must* presents a proposition as an inference; and so on” (Westmoreland 1995: 697).⁹ He further asserts an unargued for link between indirectness and weakness: “Thus the content of a sentence of the form MUST ϕ is twofold. It contains the propositional content proper, that is, it conveys the content of ϕ , which may then be added to the context in the usual fashion. Beyond that, however, the word *must* labels the content of ϕ as something known indirectly (*hence, in most cases, less certainly* [our emphasis, KvF & ASG])” (Westmoreland 1998: 79).

Some *prima facie* evidence that the speaker’s comment analysis is not entirely crazy comes from considering exchanges like the one in (8).¹⁰

⁷ See, for example, Drubig (2001) for an endorsement of Westmoreland’s proposal.

⁸ Just for the record, we take exception to the careless implication that what is inferred cannot be known. If the inference is based on logically valid inference patterns and the premises are known, the conclusion is known as well, even though “merely” inferred.

⁹ One of the present authors (von Stechow 2003) considered (and finally rejected) such a picture a few years ago.

¹⁰ We note that dialogues like this one are used by Simons in her recent work on parentheticals

(8) Q: Why isn't Louise coming to our meetings these days?

A: She might/must be too busy with her dissertation.

The crucial point is that what is proposed as the reason for Louise's absence is that she is too busy with her dissertation, not that it might or must be the case that she is too busy with her dissertation. In other words, the response in (8) offers the prejacent as the answer to the question and the epistemic modals seem to signal something like the speaker's assessment of the likelihood that this is the right answer.

If one wants to take this as evidence that modals do not contribute to the truth-conditional content of the sentence, one needs to develop an alternative semantics for them. Two possibilities are of particular interest: (i) Epistemic modals might be treated as "parentheticals", phrases that give side-remarks in a separate semantic dimension from the normal truth-conditional content. The recent treatment of such parentheticals by Potts (2005) might be thought to be adaptable to the analysis of epistemic modals. (ii) Epistemic modals might be treated as "speech act modifiers". While presenting an unmodalized sentence is interpreted as a straightforward assertion, adding an epistemic modal might indicate that a different kind of speech act (albeit with the same truth-conditional content) is performed. One might for example say that a sentence like *There might have been a mistake* expresses the speech act "I (hereby) advise you not to overlook the possibility that there has been a mistake" (cf. Swanson 2005).

The Kant/Frege approach, revived by Westmoreland and others, may not be entirely crazy but it does have one drawback: it is wrong. While locating an evidential signal in the comment dimension is (almost) on the right track, it goes wrong in treating epistemic modals as not contributing their standard quantificational force to the content dimension. The problem is that the account is inconsistent with the facts about embedding epistemic modals. While it has to be admitted that epistemic necessity modals are a bit harder to embed than their possibility counterparts, even a few such cases suffice to falsify the Kant/Frege approach. Consider, for example, yes/no questions with epistemic necessity modals. Pascal and Mordecai are playing Mastermind. After some rounds where Mordecai gives Pascal hints about the solution, Pascal asks¹¹

(Simons 2006). Here, we adapt her paradigm to the case of epistemic modals.

¹¹ The Mastermind scenario was first used in von Fintel & Gillies (2007b) for a similar purpose. Note that in many embedding environments, there is a (slight) preference for using *have*

- (9) $\left\{ \begin{array}{l} \text{Do there have to} \\ \text{Must there} \end{array} \right\}$ be two reds?

These questions clearly ask whether the evidence available entails that there are two reds, that is, the putative evidential component is directly affected by the embedding operation contrary to the expectation of the Kant/Frege approach. Similarly, while English *must* resists embedding under negation, neither *have to* nor *must*'s cognates in other languages such as German have any such compunction.¹²

- (10) a. There don't have to be two reds.
 b. *Der Code muss nicht zwei rote Stifte enthalten.*
 the code must not two red pins contain
 "There don't have to be two red pins in the code."

We conclude that the Kant/Frege account goes wrong when it claims that the epistemic modal component is not located in the content dimension.¹³

3 Kratzer and Veltman

The other way that analyses have tried to make good on the Mantra is by revising the truth-conditional entailment relation between *must*-statements and bare prejacent. There are two obvious choices: make *must* weaker or make ϕ stronger. The first path is taken by Kratzer (1991), the second by Veltman (1985).

For her Mount Toby pair, repeated here:

- (11) a. She climbed Mount Toby.
 b. She must have climbed Mount Toby.

to rather than *must*, but with no discernable difference in meaning. (It is probably good to point out that in our initial observations in (2)/(6)/(7), *have to* is interchangeable with *must*.)

¹² Languages also find it convenient to develop necessity modals that are specialized for embedding under negation such as English *need* or German *brauchen*. These have perfectly good epistemic readings.

¹³ A reviewer suggested that adherents of the Kant/Frege view could respond that our examples of embedding should be treated as involving an interaction of the modal with the question act (in (9)) and with the act of denial (in (10)). Since embedding is not restricted to just a few constructions, this would mean having to duplicate much of compositional semantics at the speech act level — not obviously a hopeless project but not one we would find attractive. Note that somehow this duplication of effort would have to extend to the analysis of negative polarity because of the facts about NPI epistemic modals mentioned in the previous footnote.

Kratzer (1991) offers this diagnosis:

In uttering (11b) rather than (11a), *I convey that I don't rely on known facts alone* [our emphasis — KvF & ASG]. I use other sources of information which are more or less reliable. These other sources may include facts concerning the normal course of events, a map, a tourist guide or hearsay. If the ordering source for the modal in (11b) is, say, a conversational background assigning to every world the set of propositions which represent the normal course of events in that world, then the proposition expressed by (11b) will not imply the proposition expressed by (11a) anymore. There are worlds w such that among all the worlds which are compatible with what we know in w , those which come closest to the normal course of events in w don't include w itself.

Thus, *must* is not simply a universal quantifier over a modal base. A *must* at w is sensitive in addition, to an ordering over possibilities — an ordering \leq_w that reflects the way things normally go at w , where what induces that ordering are generalizations like *if the rain gear is wet, it's raining* or *Billy's travel log is accurate*. The lower in the ordering a world, the more order-inducing propositions are true at that world. What *must* quantifies over isn't *all* the possibilities in B but only those which are minimal in \leq_w .¹⁴ Officially:¹⁵

Definition 1 (Weak *must* à la Kratzer (1991)).

- i. $\min(B, \leq_w) = \{v \in B : \text{there is no } u \in B \text{ such that } v \neq u \text{ and } v \leq_w u\}$
- ii. $\llbracket \text{must } \phi \rrbracket^{c,w} = 1$ iff $\min(B, \leq_w) \subseteq \llbracket \phi \rrbracket^c$

¹⁴ Occasionally it is convenient to indirectly interpret the relevant natural language constructions we will be dealing with here, associating sentences of English with sentences in some reasonable regimented language whose sentences then get assigned semantic values. We can get by with very modest means in that respect. Thus, let L be a propositional language generated from a stock of atomic sentences $A = \{p, q, r, \dots\}$, negation (\neg), and conjunction (\wedge). We will also assume that L has the one-place modal operator *must* (and its duals, when they're relevant), but are generally happy to limit our attention to such a language in which the modals only embed under negation and conjunction. For the most part, we leave the details of L in the background, and only rely on it when it matters.

¹⁵ Strictly speaking, this way of formulating the semantics of *must* is only legitimate if one makes the Limit Assumption: that there always is a set of \leq_w -minimal possibilities. Kratzer

Just because $w \in B$ it does not follow that $w \in \min(B, \leq_w)$: things might not go as expected in w . Thus *must* ϕ doesn't entail ϕ . On this analysis it is the additional evidential component of *must* — that it encodes the contribution of less certain or less direct information toward the prejacent — that makes it weak.

As we said, there is another path to weakness: make bare prejacent stronger. This is the path taken by Veltman (1985). Suppose sentences get truth values in a context c according to what partial information about the world is represented by c — the *information state* associated with context c . It doesn't matter, for our purposes, how that association gets done or what exactly an information state is. What matters is what an information state does. And what an information state does is say whether an atomic sentence is true or false or neither on the basis of it, and thereby provides a basis on which arbitrary sentences can be true or false or neither on the basis of it. And so, for present purposes, we can simply think of contexts as information states.

So fix a context c . Let B_c^+ be the set of atomic sentences of L true on the basis of it and let B_c^- be the set of atomic sentences of L false on the basis of it.¹⁶ We will want to require that $B_c^+ \cap B_c^- = \emptyset$. However, note that we will want to allow that $B_c^+ \cup B_c^- \neq A$. Gaps in our information about the world are acceptable, gluts are not.

Information states (i.e., contexts for now) admit of comparisons: one can contain at least as much information about the world as another. Take that partial ordering of information states as primitive, and constrain it as follows:

Definition 2. If c' contains at least as much information as c , $c \sqsubseteq c'$, then

- i. $B_c^+ \subseteq B_{c'}^+$
- ii. $B_c^- \subseteq B_{c'}^-$

c' (properly) extends the information in c if c' contains more information than c .

does not make that assumption, so her formulation is more complicated than the one we use here. The issue is tangential to what concerns us here: the proper treatment of the alleged weakness of *must*.

¹⁶ This is not how Veltman presents the data semantic analysis of *must*, but it is equivalent to his formulation.

Thus, whenever c' contains as much information as does c — when its partial information is at least as complete — it must be that whatever atomic sentences are settled by c are also settled by c' . Still, it is possible for c' to extend the information in c and yet as far as the *atomic* sentences go, c and c' do not differ at all. For suppose c does not rule out that $(p \wedge q)$, but that c' contains more information by ruling it out. The crucial point is that it can do this without thereby already ruling out either p or ruling out q .

This is enough to ground saying when an arbitrary sentence is true on the basis of the information in c . Two things are relevant for our purposes. First: truth conditions so grounded are partial — the fact that c might not settle whether an atomic ϕ is true or false can ramify to complex sentences which embed ϕ . Second: *must* in c quantifies over all the ways of adding to the partial picture of the world that c represents.

Definition 3 (Strong Prejacent *à la* Veltman (1985)).

$$\begin{aligned} \text{i. } \llbracket \phi_{\text{atomic}} \rrbracket^c &= \begin{cases} 1 & \text{if } \phi_{\text{atomic}} \in \mathbf{B}_c^+ \\ 0 & \text{if } \phi_{\text{atomic}} \in \mathbf{B}_c^- \end{cases} \\ \text{ii. } \llbracket \text{must } \phi \rrbracket^c &= \begin{cases} 1 & \text{if for no } c' \text{ such that } c \sqsubseteq c': \llbracket \phi \rrbracket^{c'} = 0 \\ 0 & \text{if for some } c' \text{ such that } c \sqsubseteq c': \llbracket \phi \rrbracket^{c'} = 0 \end{cases} \end{aligned}$$

The relative weakness of *must* is the result of combining the partial satisfaction conditions for bare prejacent with the quantificational force.¹⁷ A *must*-claim looks to all ways of extending the current information, checking that none of them are information states that falsify the prejacent. That is a test that can be passed even if the prejacent isn't yet settled as true. For suppose that some atomic ϕ is neither true nor false on the basis of the (direct) information associated with c , but that every way of consistently

¹⁷ Deciding the truth/falsity/whatever of atomic sentences provides a foundation for determining the truth/falsity/whatever for arbitrary sentences in a context, but it doesn't do it alone. Here are the suppressed clauses for negation and conjunction:

$$\begin{aligned} \text{iii. } \llbracket \neg \phi \rrbracket^c &= \begin{cases} 1 & \text{if } \llbracket \phi \rrbracket^c = 0 \\ 0 & \text{if } \llbracket \phi \rrbracket^c = 1 \end{cases} \\ \text{iv. } \llbracket \phi \wedge \psi \rrbracket^c &= \begin{cases} 1 & \text{if } \llbracket \phi \rrbracket^c = 1 \text{ and } \llbracket \psi \rrbracket^c = 1 \\ 0 & \text{if } \llbracket \phi \rrbracket^c = 0 \text{ or } \llbracket \psi \rrbracket^c = 0 \end{cases} \end{aligned}$$

They cover what you would expect in a sensibly constrained three-valued logic.

adding information to c results in a context c' such that ϕ is true based on c' . That is: given the information in c it will never turn out that $\neg\phi$ is true (but it will turn out one way or the other). That is clearly possible. But then given the partial information in c , $must\ \phi$ is true even though the bare prejacent ϕ is neither true nor false. Thus, $must\ \phi \not\equiv \phi$. In fact, in this set-up, $must$ is weak in the strong sense that ϕ entails $must\ \phi$: if your partial information already decides in favor of (non-modal) ϕ , then so will any consistent extension of that partial information, and so your partial information also already decides in favor of $must\ \phi$.

What is important for our purposes is that the relative weakness of $must$ is accomplished by making the satisfaction conditions for bare prejacent much stronger. Note that just like Kratzer's account, Veltman encodes a notion of evidentiality in the semantics: an information state can be seen as representing your direct evidence, and since your direct evidence can fail to satisfy ϕ even though no way of adding more evidence can falsify ϕ , it is possible that $must\ \phi$ is supported by that evidence even though ϕ isn't (yet).

We have now seen two ways of making good on the Mantra by changing the truth-conditional semantics, either by making the semantics of $must$ statements weaker (Kratzer) or by making bare prejacent stronger (Veltman). We turn to our reasons for rejecting the Mantra in either incarnation.

4 Contra The Mantra

The Mantra is that, given the basic observations, $must$ statements are weaker than bare prejacent. But that is an over-reaction to the evidentiality phenomenon and is empirically problematic.

4.1 Indirectness \neq Weakness

We first want to highlight that the Mantra gets most of its traction from moving too quickly from the basic observation about when $must$ -claims are and are not felicitous to the diagnosis that therefore $must$ must be weak. And, in particular, the confusion is to run together a feeling of *indirectness* with a feeling of *weakness*. But these are different properties. Taking the presence of the former as a reason for the latter is something that needs arguing for.

Consider, again, the characteristic scenario we began with:

- (6) [Seeing the pouring rain]
- a. It's raining.
 - b. ??It must be raining.

The *observation* is that it's plainly odd for Billy to utter the modalized claim when she has the direct information that it is raining. But the *diagnosis* that therefore the modal is weak goes beyond this. So we need an argument for it. Just noting that when Billy's information is less direct — she sees everybody coming inside with wet rain gear — she can give her weather report with *must* is not an argument. For it plainly assumes that we can conclude from the *indirectness* of Billy's evidence that what the evidence supports is *weak*. But that is just what we were supposed to be getting an argument for.

In fact, it is rare to get an argument. Instead, it is much more common to just get a glossing of the indirectness phenomenon as weakness. Thus, Karttunen (1972) says, first, that *must* ϕ makes a weaker claim than the bare prejacent ϕ . Such weak claims are appropriate, say, when ϕ follows from what is settled plus some other reasonable assumptions. But later, he says the weakness

is apparently based on some general conversational principle by which indirect knowledge — that is, knowledge based on logical inferences — is valued less highly than 'direct' knowledge that involves no reasoning.

But indirect knowledge is still knowledge, and so what follows from what is indirectly known must be true, and so there is no good sense in which *must* is weak.¹⁸

Our point is simple: weakness and indirectness are not two sides of a single coin at all. They are just different. Any arguments for a weak semantics need to be more than just reminders that *must* carries an indirect evidential

¹⁸ Similarly, Veltman (1985: p. 167) motivates his data semantic treatment by appeal to the felt indirectness associated with *must* by saying that it conveys that “the data constitute at best indirect evidence for ϕ .” He then immediately says that what is therefore needed is a semantics on which ϕ asymmetrically entails *must* ϕ , and not the other way around as it is in the standard set-up. The first part, we claim, is right and his semantics gets it right: indirect evidence for a bare prejacent ϕ is still entailing evidence for ϕ in the sense that if *must* ϕ is true on the basis of such evidence then it will eventually always turn out that ϕ . So indirectness is not the same as non-entailing. But the second part, we will argue below, isn't right: it is not right to model this indirectness as weakness.

signal.¹⁹ In any case, in what follows we will argue that there is no weakness in *must*. We first show that there are cases where it is clear that there is no weakness associated with *must*, so *must* is *not always* weak. We then move on to stronger arguments showing that *must* is never weak. These arguments will leave the Mantra in tatters. And in all cases, the indirectness of the evidence continues to be signalled. It is around the evidential signal thus that a proper treatment of epistemic *must* should be built and no amount of weakness should be impugned to it.

4.2 *Must* is Not Always Weak

Argument 4.2.1 Chris has lost her ball, but she *knows* with full certainty that it is either in Box A or B or C. She says:

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

The modal is perfectly fine. There is no good sense in which the *must* in (12) is weak.²⁰ But it continues to carry the indirect evidential signal: had Chris just opened Box C to begin with, seeing her ball inside, it would be plainly weird for her to report *It must be in C*.

¹⁹ We have not seen such arguments made explicit, although as a reviewer points out to us, one could imagine how they might go. Any time a conclusion is based on a non-trivial combination of more than one piece of evidence, possible sources of error are multiplied. So, if one has merely indirect evidence for a conclusion, one is thereby in an epistemically weaker position than if one had a single piece of direct evidence. And, then, as the reviewer writes, “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.” We are less impressed by this possible reasoning than the reviewer seems to be. First, if two pieces of evidence are known to be true, then their logical combination will not suddenly suffer from weakness just because they have been combined. Second, even if one could maintain that certain conclusions are open to epistemic worries, there is no reason to think that this would be reflected in the semantics of sentences that express those conclusions. For example, universal generalizations might be inherently more risky than existential claims, but nobody in their right mind would think that this is good reason to propose that *every* means *maybe not quite every*.

²⁰ We should note that Werner (2006: 244ff.) discusses cases such as (12) and claims that the use of *must* even there “reveals a certain hesitation” or “tentativeness”. We fail to see any indication that this is true. Admittedly, this may be a hard argument to put on solid experimental grounds. But since the literature is replete with declarations of intuitions of weakness, it seems fair to point out cases where we get a clear intuition of strength.

Argument 4.2.2 Or consider Billy. She sees people coming inside with wet rain gear. And suppose she knows *for sure* that the only possible reason for wet rain gear in these parts is rain. She can say *It must be raining* even though, *ex hypothesi*, the prejacent is entailed by the information she has. If instead she sees the pouring rain then it is plain weird for her to report it this way, even assuming that she knows that the only reason for wet rain gear in these parts is rain. Again, no weakness and persistence of the indirectness signal.

So the pattern of felicitous reports that surround Karttunen's Problem are exactly the same even if we assume that the "indirect information" involved flat-out entails the prejacent in question. When Chris is looking for her ball and knows the space of possibilities, she can say *It must be in C* when she has eliminated the other two possibilities. But she can't say that if she opens Box C and sees the ball. When Billy knows that the only reason for wet rain gear is rain, she can say *It must be raining* when she sees the wet umbrellas and slickers and galoshes. But she can't say that if she looks out the window and sees the pouring rain. Thus, the basic observation has everything to do with the directness of the speaker's information, not with the strength of the claim they make on the basis of it.

Argument 4.2.3 Another case showing that *must* can easily be used in contexts where there is no weakness attending the conclusion that the prejacent is true:

- (13) A: They said it was going to rain. I wonder whether it has started.
B: I don't think so, it was still dry when I came in 5 minutes ago.
A: Look, they're coming in with wet umbrellas. There is no doubt at all. It must be raining now.

Notice that A explicitly combines the absence of doubt with a *must*-claim. This seems to us the most natural thing in the world. Some people have suggested to us that in a way, saying *there is no doubt at all* does convey weakness — in a "the lady doth protest too much, methinks" kind of sense. This may be so, but we would urge that this not be taken as an excuse to weaken the semantics of *there is no doubt at all* to make it mean something like *there is a little bit of doubt*. That way madness lies.²¹

²¹ We come back to this point at the end. A reviewer points out to us that Grice (1989: 33-34) noted that statements of certainty lend themselves to implicatures that the speaker may not be as certain as they claim to be. Grice even uses the same Shakespeare phrase about the protesting lady as we do here (although as many others, he misquotes Shakespeare slightly).

4.3 *Must* is Never Weak

At this point, Mantraistas might consider a partial strategic retreat and admit that there are *some* cases where *must* is as strong as it ever was. But surely, all that was ever claimed by the Mantraists was that *must can sometimes* be an expression of weakness. and that the semantics needs to make space for such weak uses.²² We doubt that the authors that embraced the Mantra were intending to have their observation be about a subset of uses of epistemic *must*, but be that as it may, we will now provide evidence that in fact, *must* is *never* weak.

Argument 4.3.1 Consider the following argument:

- (14) If Carl is at the party, then Lenny must be at the party.
Carl is at the party.
So: Lenny is at the party.

The argument has this form:

- (15) If ϕ , *must* ψ
 ϕ
 $\therefore \psi$

The premises seem to entail the conclusion and not just because of the particular lexical material in (14) but because the argument form is valid. But not if the Mantra were right. For suppose the content of *must* ψ is relatively weak compared to the bare prejacent ψ . Then — no matter whether that is due to the weakness of *must* or the strength of the bare prejacent — the premises are too weak to get the conclusion.²³

²² Someone who actually makes this move is Portner (2009: p.72, fn.12), who says of an early draft of this paper that we “give examples in which a sentence with *must* carries no implication of weakness. However, Kratzer does not claim that sentences with *must* signal weakness or are invariably weaker than ones without. The only fact to be explained is that sentences with *must* are sometimes weaker, and Kratzer’s theory makes the correct predictions.”

²³ Note that the point is not that the weak *must* accounts can’t deliver *must* ϕ as the conclusion to the argument. That they presumably can. Our point here is that the bare ϕ is also a valid conclusion and that the Mantra-analyses cannot deliver that conclusion. A reviewer suggests that there might be instances of the argument form (15) that do not appear valid, especially those where a Kratzerian stereotypicality ordering source might be at play:

- (i) If Carl comes in with wet hair, it must be raining out.
Look, Carl came in with wet hair.

Argument 4.3.2 The Mantra claims that *must* is weak. That is controversial. But there are a lot of weak modals about which there is no controversy. How *must* interacts with these poses problems for the Mantra.

Suppose the Mantra is right, that *must* ϕ does not entail ϕ . Then *must* ϕ is perfectly compatible with *perhaps* $\neg\phi$. But that isn't how things sound, no matter the order of the conjuncts:

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

Sentences like the ones in (16) walk, talk, and quack like contradictions. Notice that the trouble here isn't that $\neg\phi$ *but it must be that* ϕ walks, talks, and quacks like a contradiction (though it does). That might be explainable on general pragmatic grounds: in successfully uttering $\neg\phi$ the speaker might well restrict the modal base to include only $\neg\phi$ worlds, and so the best or most normal worlds in that modal base couldn't be ϕ -worlds.²⁴ But that is not our worry. Our worry has to do with the weaker *perhaps* $\neg\phi$. Whatever else a speaker does by uttering *perhaps* $\neg\phi$, she definitely does not reduce or suggest reducing the modal base to include only $\neg\phi$ -worlds. And that is enough to cause trouble for the Mantra.

The contradictory nature of (16) makes clear that the problem with the Mantra goes deep. Our arguments do not merely show that *must*-statements

<?> So: it is raining out.

We do not share the reviewer's feeling that (i) is not valid. Whatever worries we might have are about the soundness of the first premise.

²⁴ This attempt at explaining away is like a similar attempt for dealing with Moore's paradox. Consider:

- (i) It is raining but I don't believe it.

This too has the habits, ambulatory and otherwise, of a contradiction. But, some say, that can be explained away. (See, e.g., DeRose (1991) and the references therein. It's not quite consensus, though, that the feeling of inconsistency ought to be explained away (Gillies 2001).) It *feels* like a contradiction all right, but (they say) that feeling is pragmatic. A speaker who utters (16) represents herself as knowing that it is raining when she utters the first conjunct and then says that she doesn't have the relevant belief when she utters the second. Even though the conjunction is consistent, what she represents is incompatible with what she says. (She can't know without believing.) But no such story will help with (15). For if uttering *must* ϕ "represents" that you know ϕ , then the Mantra is already in trouble. How can it represent that given (i) that knowledge is factive and (ii) the supposed weakness of *must*? (Not a rhetorical question.)

aren't *always* as weak as claimed by the Mantra. Instead, our arguments show that *must*-statements *never* have weakness as part of their meaning. Otherwise, for example, (16) should be contingently acceptable, which it isn't. It's not even close.²⁵

Argument 4.3.3 What gets this problem going in the first place is precisely the fact that *perhaps*, like *may* and *might*, is an uncontroversially weak modal. It just takes one relevant ϕ -possibility for *perhaps* / *may* / *might* ϕ to be true, and (of course) the speaker issuing such a weak modal claim doesn't have to think that the bare prejacent is true. She just has to think that the weak conditions for its truth are met. That is why — as we've argued elsewhere (von Fintel & Gillies 2007a,b, 2008) — speakers can stick to their conversational guns when they issue such claims if the prejacent turns out to be false. Thus:

- (17) a. Alex: It might be raining.
b. Billy: [Opens curtains] No it isn't. You were wrong.
c. Alex: I was not! Look, I didn't say it *was* raining. I only said it *might* be raining. Stop picking on me!

Alex doesn't have to reply like this, but she can. That holds good whether her original weak modal claim is *perhaps* or *may* or *might*. And it is not just existential modals that allow this. Take, for instance, *should* or *ought* (in their epistemic senses). These are weak necessity modals (von Fintel &

²⁵ One reviewer suggests that a possible reaction to the data in (16) is that *perhaps* is not actually an expression of weak epistemic possibility. (From the Mantra point of view, it's the dual to weak *must*, and so stronger than a pure existential.) But we're not wedded to a pure existential story about *perhaps* (though that sounds pretty good) and our argument doesn't rely on it. There *are* strong necessity epistemic modals. So pick one and take its dual (e.g., *there is a vanishingly small chance that*). It'll be horrible when paired with *must* in examples like (16), we promise. The other reviewer adduces an example found via Google:

- (i) If it's New Orleans and the novel's main characters have been dead for years but are still walking around terrorizing people, it must be an Anne Rice adventure. But it isn't [...].

(<http://www.amazon.com/City-Masks-Cree-Black-Novel/dp/1582343594>)

Clearly, we have to say that this is not straightforward conjunction of a *must*-statement with the negation of the prejacent. We agree with the reviewer that we're committed to treating this as a case of a rhetorical flourish where the second sentence retracts the statement made by the first sentence. We leave it to the reader to decide whether a semantics of *must* should take examples like (16) more seriously than we do.

Iatridou 2008). Chris is usually in her office on Tuesdays at 3. Not always (she sometimes has to go to a committee meeting then), but usually. It's Tuesday at 3 o' clock, and Alex and Billy are wondering whether Chris is free for coffee:

- (18) a. Alex: Chris oughtta/should be in her office.
b. Billy: [Opens door] No she isn't. You were wrong.
c. Alex: I was not! Look, I didn't say she *was* in her office. I only said she *ought* to be. Stop picking on me!

Again, Alex doesn't have to dig in her heels, but she can.

The common thread is weakness: the existential *might* and the weak necessity *ought* are both weak modals, and so allow Alex to distance herself from the truth of the prejacent when it turns out to be false. The point for our purposes is that *must* does not allow Alex that distance, and this is reason to think that it can't be weak:

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

Alex just can't do that. This is not puzzling from the strong *must* point of view: saying *I only said it must be raining* is as bizarre as saying *I only ate all of the cookies*. But if, as the Mantra maintains, *must* is not located at the very top of the scale of epistemic strength, one would expect *only* and *must* to combine like old friends.

Argument 4.3.4 The same underlying point can be seen from the fact that when a smidgen of weakness is called for, speakers don't regularly reach for *must*. The Hollywood crew has been hard at work setting up their rain equipment for shooting the big scene. You see people coming inside with wet umbrellas, slickers, and galoshes. You are pretty sure that rain is the only explanation — filming on the big scene doesn't start until tomorrow — but there's a twinge of doubt. What do you say? Here are some options:

- (20) a. It is raining.
b. It must be raining.
c. It is probably raining.

There is no controversy over the weakness of *probably*. There is no controversy over whether *probably* ϕ reflects that the speaker has ample credence in ϕ but leaves open the possibility that, still, $\neg\phi$. If the content of *must* really were relatively weak, we would expect something close to indifference between the *must*-claim (2ob) and its *probably* counterpart (2oc). But that is not what we see: the latter is clearly preferred to the former.

We conclude that there are serious problems with the Mantra and thus with the analyses that are designed to make good on it. Instead, we think that the basic observations that motivated the Mantra should be captured by incorporating the evidential component of *must* in a way that does not weaken its force as a strong necessity modal. That is how to solve Karttunen's Problem.

5 The Question of Direct Evidence

Our account is a combination of the standard strong semantics for *must* with an evidential component. It is the evidential component that explains the initial motivating observations. And, the standard strong semantics explains why *must* does not behave as predicted by the weakness Mantra. In many ways, the Kant/Frege approach was almost right: its main problem was that it kept the notion of epistemic necessity out of the primary semantic content. But the idea that epistemic modals bring with it an evidential signal was right on the mark. What we propose to do is to give an explicit semantics that incorporates that evidential signal and furthermore, we will attempt to characterize more fully what it means to say that something is only indirectly known — something that most current accounts of evidentiality have not attempted.

What remains to be done to fill out our proposal is two-fold: (i) we need to say what the formal status of the evidential marking of indirect inference is; and (ii) we need to say what it means to say that some proposition ϕ is known only through indirect inference. We'll take these tasks on in turn.

We have not found a language whose expression of epistemic necessity fails to carry an evidential signal of indirect inference. That is, the paradigm illustrated for English in (6)/(7) can be replicated in language after language. This should raise the suspicion that what we are dealing with should not be a stipulated, arbitrary part of the lexical meaning of epistemic necessity modals, and so it shouldn't be a lexically specified presupposition or conventional implicature. Rather, one would suspect and hope that the evidential

signal can be derived as a predictable conversational implicature that is non-detachable in Gricean terms. How could we proceed to formulate such a derivation? The obvious idea would be to say first that *must* carries no lexicalized signal of indirect inference and that it simply has the standard meaning of universally quantifying over worlds compatible with what is known. And then one would say that *must* competes with a stronger expression that is only appropriate if the prejacent is known either directly or through trustworthy reports. Then, one could derive via more or less standard quantity implicature reasoning that choosing the weaker *must* implicates that the prejacent is not known directly nor through trustworthy reports. The problem with this derivation is that there is no plausible competitor to *must* that would carry such a strong meaning. In particular, the bare assertion of the prejacent does not convey that the truth of the prejacent is known directly or through trustworthy reports. After all, it is perfectly felicitous to say that it is raining (instead of that it must be raining) upon seeing wet raingear. So, in the absence of an appropriate stronger competitor, a derivation of the indirect inference signal carried by *must* via conversational implicature reasoning cannot get off the ground.

We thus see no choice but to stipulate the evidential component of *must* in its lexical semantics and we have to leave as unsolved the mystery of why this seems to be happening with every epistemic necessity modal that we have come across.²⁶ We'd be more than happy to be shown that there is a reliable conversational derivation, but for now we will treat the evidential signal as hardwired. Now, there are at least two obvious choices of how to build the evidential signal into the semantics of *must* without depriving *must* of its strength as an epistemic necessity modal: (i) we could make it a *presupposition* of *must* that the prejacent is known only through indirect inference rather than through any kind of direct evidence or trustworthy reports; (ii) we could make that component into a *conventional implicature* à la Potts (2005).

We are convinced that the signal cannot be a conventional implicature

²⁶ Chris Potts (pc) suggested to us that perhaps there is a strong systematic pressure for languages to have a conventional signal for indirect inference and that epistemic modals are simply the perfect choice for carrying that meaning. But this wouldn't explain why *all* epistemic necessity modals (and as we'll see soon, even existential/possibility modals) carry the evidential signal. For example, why do English *must*, *have to*, *need to* all carry the signal? We leave this persistent and cross-linguistically stable pairing of epistemic modals with an indirect inference signal as a mystery. Our lexical presupposition analysis in the text is simply placeholder for the eventual solution to the mystery.

and is most likely a presupposition. One potential obstacle for a conventional implicature analysis is the feature of Potts' system that prohibits the existence of expressions that contribute both to the "at issue" dimension of meaning (the standard truth-conditional semantics) and the conventional implicature dimension. But that is precisely what we need for epistemic *must*: it contributes standard modal necessity to the "at issue" dimension and furthermore comes with the evidential signal. By Potts' logic, that evidential signal therefore cannot be a conventional implicature.

Another, probably more cogent reason to go the presupposition route comes from presupposition projection. In this connection, we actually need to refine our characterization of the indirectness. If we said that *must* comes with a presupposition that the prejacent is known only through an indirect inference, that would not quite be right yet. Recall that we need to allow for embedding epistemic necessity modals under negation as in *There don't have to be two reds*. If the epistemic modal in that sentence were to trigger the presupposition that the prejacent is known (only) through indirect inference, that would be in conflict with the assertion that results once negation applies to the modalized sentence. Instead, what we need to say is that the epistemic modal comes with a presupposition that neither the prejacent nor its negation is known through direct evidence or trustworthy reports. Then, both *There must be two reds* and expression *There don't have to be two reds* presuppose that the truth/falsity of the prejacent *There are two reds* is not directly settled by direct evidence or trustworthy evidence. Let's make this kind of talk a bit smoother. We'll call the relevant information that does not directly settle the truth/falsity of the prejacent the "privileged information". And instead of the "truth/falsity" of the prejacent, let's talk about "the question of the prejacent". So, the presupposition of epistemic modals is that the privileged information does not directly settle the question of the prejacent.

Epistemic *must* isn't special: we fully expect this presupposition of unsettledness by the privileged information to be carried by epistemic possibility modals as well. That makes our analysis uniform and principled, even though the evidential signal is posited as part of the lexical meaning of epistemic modals.²⁷ Of course, when one looks at a sentence like *It might be raining*, the

²⁷ Swanson (2008) argues that weak necessity modals like *should* do not carry exactly the same evidential signal as strong necessity modals like *must* and that thus hardwiring the signal is inescapable. While we do hardwire the signal in the absence of a workable conversational derivation, we are not yet convinced that the evidential signal is not the same (indirectness) for all epistemic modals, since the differences Swanson discusses may be due to difference

speaker — via a Gricean implicature — is heard as conveying that they are not (yet) convinced that it is raining. Therefore, the signal of unsettledness by the privileged information is “swamped” so to speak: it can’t be detected because other signals are even stronger and entail it. But our presupposition-based analysis makes a prediction: as soon as we embed a possibility modal under negation, say, we can clearly detect the signal of indirectness.

- (21) a. There aren’t two reds.
b. There can’t be two reds.

(21b) carries the signal of indirectness whereas the bare negation (21a) does not. That is: *can’t* patterns with *must*, just as we’d predict.²⁸

One last indication that we are dealing with a presupposition is provided by scenarios like the following: Alex is looking out the window, Billy figures that Alex sees the weather but actually Alex is fixated on the people in the corridor visible from there folding up wet umbrellas, slickers, and galoshes.

- (22) a. Alex: It must be raining.
b. Billy: Hey! Wait a minute. Whaddya mean, *must*? Aren’t you looking outside?

In other words, the indirectness signal passes the “Hey! Wait a minute” test for presuppositionality (von Stechow 2004).

What do we mean by the “privileged information”? It is the information that a speaker treats as given as direct trustworthy evidence, either by direct observation or through trustworthy reports. There is some natural context-dependency and vagueness here. Just about everyone will treat the direct visual observation of pouring rain as a piece of privileged directly evidenced information. But a professional epistemologist, even when on vacation in Arizona, might be tempted to say: “Well, I am getting the kind of visual input that is only consistent with rain, so it must be raining”. It is only us lucky mortals that can do without epistemic *must* and *might* at least once in a while.²⁹

at the level of asserted content only. This must remain an issue for future discussion.

²⁸ That *can’t* patterns with *must* is also (more) bad news for Mantraistas: in order to account for it, they need to either argue that epistemic possibility modals like *can* aren’t merely existential modals or they need to break duality between *must* and *can*.

²⁹ Note also that one should not equate visual evidence with privileged evidence. Observing an orchestra, auditory evidence is more privileged evidence than visual evidence. In fact, observing an orchestra through soundproof plate glass, one might

To conclude the informal sketch of our analysis: (i) epistemic modals have their standard possible worlds semantics, which includes the fact that strong epistemic necessity modals like *must* and *have to* are indeed strong; (ii) epistemic modals carry a presupposition that the question of their prejacent is not settled by the privileged information, or as we'll say most of the time below, by the "kernel". The strong semantics explains the "contra Mantra" data surveyed in Section 4. The presupposition explains the initial data that motivated the Mantra, solving Karttunen's Problem.

As we have said, the idea that epistemic modals are markers of indirect inference is not out of place with what is known about systems of evidential categories. But one thing that is rarely done in the evidentiality literature is to become formally explicit about what is meant by "indirect inference".³⁰ We will now try to fill that gap.

6 The Basic Analysis

The Mantra is wrong. We say that the basic quantificational analysis is nearly right, and so too is the comment-dimension analysis. But both are also not quite right. Epistemic modals express their usual quantificational meanings relative to a contextually determined modal base. So *must* is strong at the level of content. But such modals signal that some privileged part of the modal base, the kernel, does not directly settle the question of the prejacent either way. This signal lives somewhere outside the main at-issue content dimension—we take it (subject to the hedges we issued earlier) to be a presupposition—and is a signal about indirectness, not about weakness.

That's the basic idea. We now want to show that it can be made precise enough to explain the facts about *must*. We will sketch two (non-equivalent) implementations of that idea. The versions share a common core, but differ

have occasion to say "They must be playing Mozart". Thanks to Sebastian Löbner (p.c.) for this insight.

³⁰ There is an initial attempt at a formalization in the seminal Izvorski 1997. Izvorski proposes that the modal base for indirect evidential modals is constituted by the set of propositions $\{p: \text{speaker considers } p \text{ indirect evidence in } w\}$. But this will not work: seeing wet rain gear is indirect evidence for rain but direct evidence for the proposition that people are fully clothed. In other words, one cannot identify a set of propositions that are indirect evidence *simpliciter*. Our strategy is to start at the other end: we identify a kernel of directly given propositions. As a first approximation, any proposition that is not in that set itself but is entailed by that set is one that we have only indirect evidence for. Epistemic modals signal that the prejacent is of that nature.

at the margins. We begin with the basic analysis and then turn to the two implementations of it. Either version — indeed, *any* plausible way of implementing our basic analysis — can solve Karttunen’s Problem.

Here is what our implementations will not do: they will not say that the evidential signal associated with *must* ϕ is a matter of labeling the prejacent with a suitable tag like INDIRECT or INFERRED. First, because we don’t see what’s gained by changing to smallcaps. Second, because such a story is wildly unconstrained — we can think of too many labels that we can write down, but that never find their way to being carried as an evidential signal. The tag INFERRED ON A TUESDAY FROM WHAT MAURY SAID WHILE SITTING AT THE LUNCH COUNTER AFTER SENDING BACK THE SOUP THAT WAS SUPPOSED TO BE A BISQUE BUT WAS REALLY JUST CREAM OF TOMATO is perfectly good *qua* label, but not perfectly good *qua* evidential signal. Instead, we want the evidential signal to emerge from combining the basic ingredients of the semantics of modals.

We begin with just the materials that the basic modal semantics provides: a world and a contextually supplied modal base at that world. The modal base represents the information compatible with what is known at the world in the context. When it doesn’t cause confusion, we can simply take modal bases to be sets of contextually relevant worlds. Our proposal is that not all information that our modals quantify over is created equal: some of it is privileged information. So we also need, in addition to modal bases, *kernels*. These represent what information is *direct* information in the context — or direct *enough* in the context, since what counts as direct may well depend on context.

Let’s represent the direct information in a context with a (finite, non-logically closed) set of propositions — those propositions encoding the direct information. The modal base determined by a kernel of information is simply the image thereof:

Definition 4 (Kernels and Bases). K is a kernel for B_K , B_K is determined by the kernel K , only if:

- i. K is a set of propositions (if $P \in K$ then $P \subseteq W$)
- ii. $B_K = \bigcap K$

You can’t have direct information that P unless it is the case that P . So for a modal uttered at w , with respect to a kernel K , we know that $w \in \bigcap K$. So our modal bases will be reflexive. Given this set-up, there is a special

case: let $K_o = \emptyset$ be the minimal kernel. It follows that $B_{K_o} = W$ — if you have absolutely no direct enough information, you have no information full-stop.³¹

By treating the evidential signal as a presupposition we can then impose its satisfaction as a definedness constraint on the assignment of truth values in a context at a world. Doing that gives us what is common to our two versions of implementing our proposal. Here is the entry for the strong necessity modal *must*:

Definition 5 (Strong *must* + Evidentiality). Fix a c -relevant kernel K :

- i. $\llbracket \textit{must } \phi \rrbracket^{c,w}$ is defined only if K does not directly settle $\llbracket \phi \rrbracket^c$
- ii. $\llbracket \textit{must } \phi \rrbracket^{c,w} = 1$ if $B_K \subseteq \llbracket \phi \rrbracket^c$

As we will see, the two different implementations differ in what they say about how a kernel directly settles a proposition.

The basic intuition is that K can *fail* to directly settle whether P even though K *entails* whether P ; epistemic modals carry an evidential signal that exploits that gap. Even before we make that intuition more precise we can see how a story that captures it can help. Take the original example:

- (6) [Billy seeing the pouring rain]
 - a. It's raining.
 - b. ??It must be raining.

Here Billy has direct information — intuitively, the contextually supplied kernel directly settles — that it is raining. The kernel settles the prejacent and this conflicts with *must*'s evidential signal. That's why (6b) is bad. It's

³¹ Two points about kernels. First, kernels are, of course, simply classic Kratzerian modal bases (Kratzer 1977): such a base at a world is a finite, non-closed set of propositions representing what is known at the world in question. She had her reasons for wanting modal bases to have extra structure over and above the image of such a base — as a result modals, on that picture, don't end up acting simply as quantifiers over a domain of possibilities. Our story uses the same added structure but puts it to new use in a way that preserves the simple quantificational picture of the modals. So we think it is useful to separate how we represent the direct information (our kernel) from how we represent the domain of quantification (our modal base). Second, everything we say in the text assumes that the only information is either direct (enough) or follows from what is direct. That is a substantive claim that simplifies the presentation. But 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 of (as we do in the text) taking it to be W .

not that Billy says something weak when she says *It must be raining*. It's that the modal signals her information isn't direct when it is. But when Billy instead has indirect but no less conclusive information about the weather, she can report the weather with *must* if she wants:

- (7) [Billy seeing wet rain gear and knowing rain is the only cause]
- a. It's raining.
 - b. It must be raining.

Here Billy's direct information is that the people coming inside have wet umbrellas, slickers, and galoshes and that rain is the only cause. Now her choice isn't forced: (22b) presupposes that the kernel doesn't directly settle whether it is raining. But the basic intuition is that a kernel can fail to directly settle whether P even though it entails whether P . So there is room here for (22b)'s presupposition to be met and Billy's utterance of it to be felicitous (and true). She doesn't *have* to say (22b), of course. The bare (22a) would be just fine, too. That is a fact that is a little uncomfortable for the Mantra but our story takes in perfect stride.

Our claim isn't that *must* carries this extra evidential component over and above its quantificational oomph. Our claim is that epistemic modals carry this signal. In particular, *must*'s dual, *might/can* carries it as well:

Definition 6 (Weak *might* + Evidentiality). Fix a c -relevant kernel K :

- i. $\llbracket \textit{might/can } \phi \rrbracket^{c,w}$ is defined only if K does not directly settle $\llbracket \phi \rrbracket^c$
- ii. $\llbracket \textit{might/can } \phi \rrbracket^{c,w} = 1$ if $B_K \cap \llbracket \phi \rrbracket^c \neq \emptyset$

As we said, since *might/can* is an existential modal, this signal can be hard to spot. If your direct information settles whether it is raining, then you really shouldn't be saying *It might be raining*. What you say is either misleadingly weak or just plain false. But notice how easy it would be to spot the evidential signal when the existential modal is embedded under negation (e.g., in *can't*) — if, as we say, the signal is a presupposition. For then we would expect the presupposition to project to the entire claim. And that's just what we find. When on her trip to the desert southwest, Billy can report with (23a) but not (23b):

- (23) [Billy seeing brilliant sunshine]
- a. It's not raining.

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

Later, Billy has to go indoors to attend the conference (in windowless room to make sure people pay attention to the speakers). She sees people coming inside folding their sunglasses, hats, and parasols and knows that brilliant sunshine is the only explanation. Now her report is no longer constrained:

(24) [Billy seeing people with sungear, knowing sunshine is the only cause]

- a. It's not raining.
- b. It can't be raining.

So our claim that epistemic modals generally carry this evidential signal of indirectness, plus our hypothesis about how that signal is carried, naturally predicts that *can't* patterns with *must*.³² Just as with the observation about *must*, we can see how a story that exploits the gap between what a kernel directly settles and what it entails and contradicts — even without yet giving a precise way of making that distinction — can explain this phenomenon.

7 Two Implementations

The basic analysis (Definition 5 for *must* and Definition 6 for the dual possibility modals) makes an appeal to whether a kernel fails to directly settle the question of the prejacent. Even before making that idea precise, we can get some way towards an explanation of the observation that motivates the mantra. But that does depend on it making sense to distinguish between what the direct (enough) information in a context directly settles and what that information entails. If there isn't a sensible way of saying what the gap is between what is directly settled by a kernel *K* and what follows from *K*, then the kind of explanation we want to give won't hold up. We now want to consider two (non-equivalent) ways of saying how kernels directly settle issues that live up to that requirement. Either of them can be paired with the basic analysis.

So we need a sensible way of modeling the difference between what is

³² How the evidential-signal-as-presupposition hypothesis fares generally is, as we said, another topic for another day — we don't, for example, make any (special) claims about how this presupposition projects from under attitude verbs. While we're at it: we also make no interesting claims about the semantics of attitude verbs, and so we are happy enough to inherit whatever virtues and vices possible worlds stories of them have.

directly settled by a set of propositions and what follows from that set of propositions. Here we have good news and bad news. The good news is this isn't a new problem. The bad news is this is (pretty much) the problem of logical omniscience — the problem of distinguishing between what an agent *explicitly* knows and what she knows in virtue of it following from what she explicitly knows. But the *other* good news is that we therefore know the rough menu of options for modeling what it means for a kernel to directly settle whether P .³³

7.1 First Implementation

The first way of implementing our basic analysis takes a direct path: whether some question is directly settled boils down to whether there is an independent bit of direct information that answers the question. Given a kernel K , we simply say that the question whether P is a directly settled issue with respect to K just in case either P is entailed or contradicted by one of the pieces of direct information explicitly given by the context. That is:

Implementation 1 (Explicit Representation). K directly settles whether P iff either $X \subseteq P$ or $X \cap P = \emptyset$ for some $X \in K$.

This strategy, as we said, has close relatives: belief base models of belief revision (Hansson 1999) and Montague-Scott models of epistemic logic. Opting for it would also reduce the distance between our story and Kratzer's story about non-logically closed modal bases. This leaves room for a gap between what is directly settled by a kernel and what follows from a kernel. For example: suppose we have the situation represented in Figure 2. Here the kernel contains two propositions, P and Q . Neither of these on their own either entails or contradicts the proposition R . Still, the modal base this kernel determines does settle whether R since $P \cap Q \subseteq (W \setminus R)$.

³³ As we'll see, we could: (i) opt for some method of explicit representation of what is directly settled (e.g., Hansson's (1999) belief base models of belief revision, Montague-Scott models of epistemic logic); (ii) opt for a partial satisfaction accounts, based more or less on the possible worlds semantics for intuitionistic logic (e.g., Veltman's (1985) data semantics, and so-called "non-standard models" of epistemic logic); or (iii) opt for an account that distinguishes uncertainty or ignorance about which points obtain from ignorance or uncertainty about what points there are — partition models are good for that (Halpern & Rêgo 2005; Heifetz, Meier & Schipper 2006; Li 2006). For a survey of models that have the problem of logical omniscience at the forefront see Fagin, Halpern, Moses & Vardi (1995).




Figure 2 Explicit Representation: An Example

Pairing this way of modeling directness with Definition 5 explains the basic observation about *must*. Let's adopt some abbreviations:

- (25) a. $P = \llbracket \text{no wet rain gear} \rrbracket$
 b. $Q = \llbracket \text{it's raining} \rrbracket$

Now suppose the context is one in which Billy is indoors and sees people coming in with wet umbrellas, slickers, and galoshes (and knows rain is the only explanation) — a context, in other words, that determines a kernel representing that there is wet rain gear: $K = \{P \cup Q, W \setminus P\}$. In that context, Billy can say (7b). The modalized report is just fine. That's because there's no single proposition in K that entails or contradicts Q ; no such proposition is explicitly given by the context. And so (by the lights of Implementation 1) that question isn't directly settled by K . Thus our basic analysis says that the presupposition of *It must be raining* is met. And since $B_K \subseteq Q$, it is also true.

If, instead, Billy is in a context in which she has direct information that it is raining, then the package deal of Definition 5 + Implementation 1 explains why she can't say (6b). In such a context we have a kernel representing that it is raining: $K = \{P \cup Q, Q\}$. Obviously, $Q \subseteq Q$ so there is a proposition in K that either entails or contradicts Q and so (by the lights of Implementation 1) Q is directly settled by K . Thus our basic analysis says that the presupposition of *It must be raining* isn't met. That's why it's odd for Billy to say it, even though $B_K \subseteq Q$.

All of this is reflected in Figure 3. If Billy has direct information that people are coming in with wet umbrellas, slickers, and galoshes ($W \setminus P$) and the direct information that rain gear gets wet only if it is raining ($P \cup Q$), then the situation is as in Figure 3(a). No piece of her direct information




Figure 3 Explicit Representation: Billy's Weather Reports

settles the question about the weather (each such proposition overlaps with, but is not included in, the Q region). So the presupposition of *it must be raining* — that $\llbracket it\ is\ raining \rrbracket^c$ isn't settled by K — is met. Still, the two pieces of direct information together entail Q . So *it must be raining* is true.

If, on the other hand, Billy just looks out the window and sees the pouring rain, then things are as Figure 3(b) paints them. As before, she knows that wet rain gear invariably means that it is raining ($P \cup Q$). But she doesn't see the wet rain gear. She sees the pouring rain (Q). Now she has a piece of direct information that settles the question about the weather — the Q region completely coincides with the Q region, after all. The evidential signal of *it must be raining* is the presupposition that K doesn't settle whether Q . Since it clearly does, Billy can't issue her weather report with the modal even though $B_K \subseteq Q$.

If offered Definition 5 + Implementation 1 as a package-deal, take it or leave it, we take it. But the implementation is an optional extra not forced on us by the analysis. So we want to now sketch an alternative, and non-equivalent, implementation that can be paired with our basic analysis with the same effect.

7.2 Second Implementation

Our intuitive gloss says that epistemic modals presuppose that the direct information in a context doesn't settle the question of the prejacent. So when Billy's direct information is about whether it is raining, the modalized

weather report is ruled out. The second implementation begins here, saying what it takes for a kernel to be about a proposition.

A kernel K determines a modal base B_K but it also determines a set of issues — a subject matter. Before saying how it can do that, we want to say something about subject matters. First: they are equivalence relations over W . That is equivalent to saying that they are partitions, the cells of which are the equivalence classes in the relation. Either option will do. A proposition P is about a subject matter just in case any two worlds that live in the same equivalence class induced by the subject matter do not differ over whether or not P . No wonder we glossed the contribution of kernels as answers to the question *What's your direct (enough) information?* Our story here (plus or minus just a bit) is the standard partitioning semantics for questions (Groenendijk & Stokhof 1984, 1997), and that in turn is equivalent to a story about aboutness (Lewis 1988). And second: not all subject matters are created equal. Some are refinements of others. Refining is a matter of distinguishing possibilities, not ruling them out. To refine a prior partition along the boundaries of some proposition P you keep only the pairs of worlds from it that agree on whether P . When you are done, you will have a new subject matter — one in which P is an issue. Of course, some refinements are trivial refinements: the issue to be raised is already raised. That is how we tell what the issues are.

Collecting all of this up:

Definition 7 (Subject Matters, Refinement, and Issues). Let P be a proposition and S a subject matter.

- i. A subject matter S is an equivalence relation on W
- ii. $S[P] = \{\langle w, v \rangle \in S : w \in P \text{ iff } v \in P\}$
- iii. P is an issue in S iff $S[P] = S$

Again there is a special case: S_o is the universal relation over W : no issues have been explicitly raised in it because no possibilities have been distinguished. Lewisian aboutness says that a proposition P is about a subject matter S iff all the worlds in each of the S -equivalence classes agree on P . It is easy to prove that P is an issue in S iff P is about S .

The idea for this implementation is that the direct information in a context raises a set of issues. An epistemic modal claim presupposes that its prejacent isn't among those issues: it signals that the direct information in

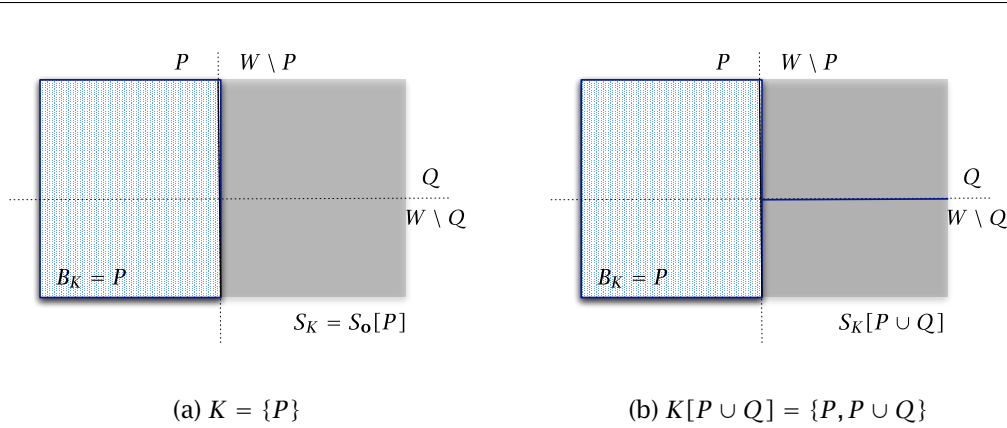


Figure 4 Settling By Partitions: An Example

the context doesn't settle the question in the relevant sense. So let K be a kernel. It determines a modal base B_K and a subject matter and thereby a set of issues:

Implementation 2 (Settling By Partitions). Let $K = \{P_1, \dots, P_n\}$. Then:

- i. $S_K = S_o[P_1] \dots [P_n]$
- ii. K directly settles whether P iff P is an issue in S_K

When it won't cause too much confusion, we sometimes write $K[P]$ for the kernel K plus the new piece of direct information that P . All of this is equivalent to saying that a kernel K partitions W along the P_i -boundaries, giving us a kernel-partition with cells K_1, \dots, K_n . Those cells are just the S_K -equivalence classes. There is nothing to choose between talking in terms of sets of issues, subject matters, and partitions. So we don't choose.

Here is a simple example (Figure 4). Suppose we are in a context in which the only direct information is that P . Then the kernel divides things up as in Figure 4(a). The dashed lines are the boundaries between propositions and their complements, the solid lines mark boundaries between equivalence classes, and the modal base is the shaded region. Compare that to a context which adds the new bit of direct information that $P \cup Q$. The kernel $K[P \cup Q]$ divides things up further as in Figure 4(b). That is because there are worlds that agree on whether P but disagree on whether $P \cup Q$. The partition S_K didn't distinguish such worlds but $S_K[P \cup Q]$ does. So these kernels have different subject matters even though they determine the same modal base.




Figure 5 Settling by Partitions: Billy's Weather Reports

Thus there is a gap between what follows from a kernel and what is directly settled by it.

Pairing this way of modeling directness with our basic analysis (Definition 5) also solves Karttunen's Problem. Consider, first, a context in which Billy's direct enough information is that people are coming in with wet umbrellas, slickers, and galoshes ($W \setminus P$) and that rain is the only explanation ($P \cup Q$). So $K = \{P \cup Q, W \setminus P\}$. In this context, Billy can say (7b). The modalized report is just fine. That's because *it must be raining* presupposes that Q isn't settled by the direct information in the context. And there are pairs of worlds in S_K that agree on $P \cup Q$ and agree on whether $W \setminus P$ but differ on Q . So $S_K[Q] \neq S_K$. Whether or not Q isn't what Billy's direct information is about and so (by the lights of Implementation 2) that information doesn't settle whether Q . That is what Billy's use of the modal report signals. Thus our basic analysis says that the presupposition of *It must be raining* is met. And since $B_K \subseteq Q$ it is also true.

If, instead, Billy is in a context in which she has direct information that it is raining, then the package deal of Definition 5 + Implementation 2 explains why she can't say (6b). In such a context we have a kernel representing that it is raining: $K = \{P \cup Q, Q\}$. Obviously, Q is an issue in S_K since once we've divided up the space along the Q -boundaries in S_K doing it again won't do anything. So (by the lights of Implementation 2) Q is directly settled by K . Thus our basic analysis says that the presupposition of *It must be raining* isn't met. That's why it's odd for Billy to say it, even though $B_K \subseteq Q$.

All of this is reflected in Figure 5. If Billy has direct information that

people are coming in with wet umbrellas, slickers, and galoshes ($W \setminus P$) and the direct information that rain gear gets wet only if it is raining ($P \cup Q$), then the situation is as in Figure 5(a). Her direct information doesn't settle the question about the weather. That is because there are worlds — we've drawn a circle showing where you might find some — that agree on $W \setminus P$ and on $P \cup Q$ but disagree on whether or not Q . Those worlds aren't distinguished by this kernel: the direct information it represents isn't about Q . So the presupposition of *it must be raining* — that K doesn't settle whether Q — is met. Still, $B_K \subseteq Q$ and so the direct information entails Q . So her modalized weather report is true.

If, on the other hand, Billy just looks out the window and sees the pouring rain, then things are as Figure 5(a) paints them. As before, she knows that wet rain gear invariably means that it is raining ($P \cup Q$). But she doesn't see the wet rain gear. She sees the pouring rain (Q). Now her direct information is about Q . That is because $S_K[Q] = S_K$: there are no worlds that agree on $P \cup Q$ and on Q but disagree on Q . Once we have divided things up along the Q boundaries, doing it again won't distinguish any more possibilities. The evidential signal of *it must be raining* is the presupposition that K doesn't settle whether Q . Since it clearly does, Billy can't issue her weather report with the modal even though $B_K \subseteq Q$.

So we have two ways of implementing our basic analysis. They both do what is required to provide a solution to Karttunen's Problem. And they both predict at no extra cost the same distribution of facts about *can't*. But the two implementations are not equivalent. Even though our interest here isn't to choose between these ways of implementing our basic analysis, we do want to point out some boundary cases where they differ.³⁴

According to the explicit representation strategy (Implementation 1) if $P \cap Q$ is direct information in a context then P is also directly settled. That's simply because $P \cap Q \subseteq P$. Similarly: if P is direct information in a context then $P \cup Q$ is directly settled. That's simply because $P \subseteq P \cup Q$. The partition strategy (Implementation 2) disagrees on both counts. For suppose $K = \{P \cap Q\}$. There are worlds that agree on whether $P \cap Q$ but disagree

³⁴ But here is another spot where they agree: $P \cup Q$ can be directly settled even though neither P nor Q is directly settled. Both implementations allow for this possibility. We think that is in their favor and why we offered them instead of an analysis based on partial satisfaction — for example, Veltman's (1985) data semantics or so-called "non-standard models" of epistemic logic (Fagin et al. 1995). Stories like that tend to predict that disjunctive propositions can only be settled if one or the other disjunct is.

over whether P . Take, for instance, a world in $P \cap (W \setminus Q)$ and a world in $(W \setminus P) \cap Q$. These worlds agree on $P \cap Q$, both delivering the verdict *false*, but disagree on whether P . So $S_K \neq S_{K[P]}$. Similarly: suppose we have a kernel $K' = \{P\}$. There are worlds that agree on P but differ on $P \cup Q$ — in fact, we saw that already in our earlier example in Figure 4. So $S_{K'} \neq S_{K'[P \cup Q]}$.

When paired with our basic analysis, this difference can make for a difference in predictions about when *must* is felicitous. But the facts here are delicate because the cases are boundary cases. An example: suppose Billy has the direct information that Alex painted a particular still life with apple and bananas. Can Billy say *She must have painted bananas*? Or suppose Billy is at the zoo, standing in front of the big gazelle-plus-antelope enclosure. She sees all the animals off in the distance toward the other end. Later can she report *I must have seen gazelles*? The package deal of Definition 5 + Implementation 2 predicts that she can. But by judicious selection of what counts as direct enough information in the context so can the package deal of our Definition 5 + Implementation 1. For instance, perhaps Billy didn't have the direct information that there were gazelles and antelopes in the enclosure ($P \cap Q$); instead perhaps the contextually relevant kernel represents the direct information that either the enclosure was empty or there were gazelles and antelopes in it ($R \cup (P \cap Q)$) together with the information that the enclosure wasn't empty ($W \setminus R$). In that case, one can get the prediction that the *must* is just fine.

But, as we said, our interest here is in offering implementations of our basic analysis not in deciding between them. That is something we'll leave for another day. Today's task was to set the story right about Karttunen's Problem. What is settled need not be directly settled: there is a gap between the direct information in a context and what follows from that information. Our modal talk exploits that. But to exploit that we need not — indeed, ought not — think that being indirectly settled is any less settled. That is the mistake of the Mantra, and — no matter the implementation — we have steered clear of it.³⁵

³⁵ This kind of explanation is similar in spirit, but not at all in execution, to the explanation in Stone (1994). What does the work in Stone's analysis is that *must* refers to an argument that conclusively settles the prejacent. (He then formalizes that by relying on a particular implementation of support between arguments from the defeasible reasoning literature.)

8 Wrap-Up and Onward and Beyond

We have argued that the Mantra that epistemic *must* is a marker of weakness is an overreaction to a misdiagnosis of the much more interesting fact that epistemic *must* is an evidential marker signalling an indirect inference. Speakers who say *must* ϕ are just as strongly committed to the prejacent as those who assert ϕ by itself. Of course, there are prejacent for which intuitively direct evidence is more convincing evidence than indirect inferential evidence. So, a speaker who chooses nevertheless to use the strong *must* ϕ incurs a higher degree of risk. So, we may judge that in many cases, *must* ϕ is more likely to be false than ϕ by itself would have been if there had been direct evidence for the prejacent. But a sentence being more likely to be false than another is far from an argument that it is weaker! As an illustration, consider

(26) I must be hungry.

Usually, we have direct perception of our internal state of hunger. (26) is unusual in that it signals that while the speaker is committed to being hungry, this is based on an indirect inference (the kernel does not directly settle whether the speaker is hungry). But by choosing *must*, a strong necessity modal, the speaker nevertheless fully commits herself to the inference. If she had wanted to be more tentative, there would have been other options: *I'm probably hungry*, etc.

This brings us back to the possibility that the strength of *must* claims may be exploited in a “the lady protests too much” kind of way (discussed earlier in Section 4.2). Just as saying *I have no doubt that he will be here very soon* may sometimes indicate the presence of at least a smidgen of doubt, saying *it must be raining* may indicate the presence of tentativeness. But just as it would be insane to attribute this kind of weakness to the semantics of *I have no doubt*, we have argued that no weakness should be built into the semantics of *must*.

Before we put an exclamation point at the end of this paper, we would like to point out just a couple of the many avenues of further exploration that we see in front of us. (i) We need to find out whether it is indeed impossible to derive the indirectness signal from conversational principles rather than hardwiring it into the semantics of epistemic modals. (ii) The evidentiality literature has rarely gone beyond labelling strategies to treat the semantic differences between various evidential categories. An obvious extension of

our work would explore whether the tool of structured modal bases can help to give a principled account of the variety of evidential systems found cross-linguistically. (iii) There are intriguing parallels and differences between epistemic *must* and expressions of clarity (Barker 2009; Bronnikov 2008) that should be explored.

Conclusion: *Must* is strong!

A Appendix: The Standard Prediction

Let us quickly rehearse the prediction of the standard quantificational analysis that *must* gives rise to stronger statements than the bare prejacent. Epistemic modals are quantifiers over worlds supplied by an epistemic modal base: *might* is an existential quantifier, and *must* its universal dual, over the possibilities compatible with “what is known”.

More precisely: a modal base B_c (supplied by a context c) determines a set of worlds, those compatible with the c -relevant information. Different c 's might call for different groups of knowers, or might call for different standards whereby the knowers know. Thereof we shall not speculate here.

The standard view officially takes (epistemic) modal bases to be functions from worlds to sets of worlds compatible with the c -relevant information at the world in question. We will insist that such functions are reflexive and euclidean — else they could not model possibilities compatible with what is *known*, in whatever context by whatever standards. Suppose f_c is such a function. Then:

- (27) a. for any w : $w \in f_c(w)$ (Reflexiveness)
 b. for any w, v : $v \in f_c(w)$ implies $f_c(w) \subseteq f_c(v)$ (Euclideaness)

That straightaway implies that such sets are *closed*: that if $v \in f_c(w)$ then $f_c(v) = f_c(w)$. Thus, once we settle on a context c , the set of worlds quantified over in c by the epistemic modals we'll be interested in do not vary between worlds compatible with c . So we can simply take a modal base in c to be a set of worlds B_c compatible with the c -relevant information, and rest assured that if a modal is evaluated at w in c then $w \in B_c$:

Definition 8 (Modal Bases). B_c is a modal base for a modal at w in c only if:

- i. $B_c = \{w: w \text{ is compatible with the } c\text{-relevant information}\}$

ii. $w \in B_c$

When the context is clear, we omit the subscript and just write ‘ B ’ for the modal base.³⁶

It is reflexivity of the function determining B that guarantees that the bare prejacent ϕ is (asymmetrically) entailed by *must* ϕ . The reason is simple: *must* in c at w is a universal quantifier over B and we know that $w \in B$. That is:

Definition 9 (Strong *must*). $\llbracket \text{must } \phi \rrbracket^{c,w} = 1$ iff $B \subseteq \llbracket \phi \rrbracket^c$.

So, if the *must*-claim is true in c at w , then $B \subseteq \llbracket \phi \rrbracket^c$. Since $w \in B$, it follows that $w \in \llbracket \phi \rrbracket^c$ — that is, ϕ is true at w . Thus, if *must* ϕ is true in c at w then so must be ϕ . So *must* ϕ is stronger than the bare prejacent ϕ .

References

- Aikhenvald, Alexandra Y. 2004. *Evidentiality*. Oxford: Oxford University Press.
- Barker, Chris. 2009. Clarity and the grammar of skepticism. *Mind and Language* 24(3). 253–273. doi:10.1111/j.1468-0017.2009.01362.x.
- Blain, Eleanor M. & Rose-Marie Déchaine. 2007. Evidential types: Evidence from Cree dialects. *International Journal of American Linguistics* 73(3). 257–291. doi:10.1086/521728.
- Bronnikov, George. 2008. The paradox of clarity: Defending the missing inference theory. *Proceedings of Semantics and Linguistic Theory* 18. doi:1813/13031.
- Davis, Christopher, Christopher Potts & Margaret Speas. 2007. The pragmatic values of evidential sentences. *Proceedings of Semantics and Linguistic Theory* 17. doi:1813/11294.
- DeRose, Keith. 1991. Epistemic possibilities. *The Philosophical Review* 100(4). 581–605. doi:10.2307/2185175.
- Drubig, Hans Bernhard. 2001. On the syntactic form of epistemic modality. Ms, Universität Tübingen. <http://www.sfb441.uni-tuebingen.de/b2/papers/DrubigModality.pdf>.

³⁶ Moreover: since dependence on a modal base is the only context dependence at issue here, we can distinguish different contexts c and c' by noting the different modal bases B and B' associated with them. We will let style be our guide.

- Fagin, Ronald, Joseph Y. Halpern, Yoram Moses & Moshe Y. Vardi. 1995. *Reasoning about knowledge*. MIT Press.
- Faller, Martina. 2002. *Semantics and pragmatics of evidentials in Cuzco Quechua*: Stanford University dissertation. <http://personalpages.manchester.ac.uk/staff/Martina.T.Faller/documents/Thesis.pdf>.
- von Fintel, Kai. 2003. Epistemic modals and conditionals revisited. Handout from a colloquium talk at UMass Amherst.
- von Fintel, Kai. 2004. Would you believe it? The king of France is back! Presuppositions and truth-value intuitions. In Marga Reimer & Anne Bezuidenhout (eds.), *Descriptions and beyond*, 315–341. Oxford University Press.
- von Fintel, Kai & Anthony S. Gillies. 2007a. *Might* made right. To appear in a volume on epistemic modality, edited by Andy Egan and Brian Weatherson, Oxford University Press. <http://mit.edu/fintel/fintel-gillies-2007-mm.pdf>.
- von Fintel, Kai & Anthony S. Gillies. 2007b. An opinionated guide to epistemic modality. In Tamar Szabó Gendler & John Hawthorne (eds.), *Oxford studies in epistemology: Volume 2*, 32–62. Oxford University Press. <http://mit.edu/fintel/fintel-gillies-2007-ose2.pdf>.
- von Fintel, Kai & Anthony S. Gillies. 2008. CIA leaks. *The Philosophical Review* 117(1). 77–98. doi:10.1215/00318108-2007-025.
- von Fintel, Kai & Sabine Iatridou. 2008. How to say *ought* in Foreign: The composition of weak necessity modals. In Jacqueline Guéron & Jacqueline Lecarme (eds.), *Time and modality* (Studies in Natural Language and Linguistic Theory 75), 115–141. Springer. doi:10.1007/978-1-4020-8354-9.
- Frege, Gottlob. 1879. *Begriffsschrift, eine der arithmetischen nachgebildete formelsprache des reinen denkens*. Halle: L. Nebert. Translation by Stefan Bauer-Mengelberg, from Jean van Heijenoort (ed.) *From Frege to Gödel: A Source Book in Mathematical Logic, 1879-1931*, Cambridge, MA: Harvard University Press. 1967.
- Gillies, Anthony S. 2001. A new solution to Moore's paradox. *Philosophical Studies* 105(3). 237–250. doi:10.1023/A:1010361708803.
- Grice, Paul. 1989. *Studies in the way of words*. Cambridge, MA: Harvard University Press.
- Groenendijk, Jeroen & Martin Stokhof. 1984. *Studies on the semantics of questions and the pragmatics of answers*: Universiteit van Amsterdam dissertation. <http://dare.uva.nl/en/record/123669>.
- Groenendijk, Jeroen & Martin Stokhof. 1997. Questions. In Johan van Benthem

- & Alice ter Meulen (eds.), *Handbook of logic and language*, 1055-1124. Elsevier.
- Groenendijk, Jeroen A.G. & Martin J.B. Stokhof. 1975. Modality and conversational information. *Theoretical Linguistics* 2(1/2). 61-112.
- de Haan, Ferdinand. 2001. The place of inference within the evidential system. *International Journal of American Linguistics* 67(2). 193-219. doi:10.1086/466455.
- Halpern, Joseph Y. & Leandro Chaves Rêgo. 2005. Interactive unawareness revisited. In *Tark '05: Proceedings of the 10th conference on theoretical aspects of rationality and knowledge*, 78-91. Singapore, Singapore: National University of Singapore.
- Hansson, Sven Ove. 1999. *A textbook of belief dynamics: Theory change and database updating* (Applied Logic Series 11). Dordrecht: Kluwer Academic Publishers.
- Heifetz, Aviad, Martin Meier & Burkhard C. Schipper. 2006. Interactive unawareness. *Journal of Economic Theory* 130(1). 78-94. doi:10.1016/j.jet.2005.02.007.
- Izvorski, Roumyana. 1997. The present perfect as an epistemic modal. *Proceedings of Semantics and Linguistic Theory* 7. 222-239. <http://www-rcf.usc.edu/~pancheva/evidentialperfect.pdf>.
- Kant, Immanuel. 1781. *Critik der reinen Vernunft*. Riga: Johann Friedrich Hartknoch.
- Karttunen, Lauri. 1972. *Possible and must*. In J. Kimball (ed.), *Syntax and semantics*, vol. 1, 1-20. New York: Academic Press.
- Kratzer, Angelika. 1977. What *must* and *can* must and can mean. *Linguistics and Philosophy* 1(3). 337-355. doi:10.1007/BF00353453.
- Kratzer, Angelika. 1981. The notional category of modality. In H. J. Eikmeyer & H. Rieser (eds.), *Words, worlds, and contexts: New approaches in word semantics*, 38-74. Berlin: de Gruyter.
- Kratzer, Angelika. 1991. Modality. In Arnim von Stechow & Dieter Wunderlich (eds.), *Semantics: An international handbook of contemporary research*, 639-650. Berlin: de Gruyter.
- Lewis, David. 1988. Relevant implication. *Theoria* 54(3). 161-174. doi:10.1111/j.1755-2567.1988.tb00716.x.
- Li, Jing. 2006. Information structures with unawareness. Ms, University of Pennsylvania. <http://www.econ.upenn.edu/~jili/research/unaware.pdf>.
- Lyons, John. 1977. *Semantics*. Cambridge: Cambridge University Press.
- Matthewson, Lisa, Hotze Rullmann & Henry Davis. 2007. Evidentials as epis-

- temic modals: Evidence from St'át'imcets. *Linguistic Variation Yearbook* 7(1). 201-254.
- McCready, Eric & Norry Ogata. 2007. Evidentiality, modality and probability. *Linguistics and Philosophy* 30(2). 147-206. doi:10.1007/s10988-007-9017-7.
- Papafragou, Anna. 2006. Epistemic modality and truth conditions. *Lingua* 116(10). 1688-1702. doi:10.1016/j.lingua.2005.05.009.
- Portner, Paul. 2009. *Modality*. Oxford University Press.
- Potts, Christopher. 2005. *The logic of conventional implicatures*. Oxford University Press.
- Rooryck, Johan. 2001a. Evidentiality, part 1. *GLOT International* 5. 125-133.
- Rooryck, Johan. 2001b. Evidentiality, part 2. *GLOT International* 5. 161-168.
- Simons, Mandy. 2006. Observations on embedding verbs, evidentiality, and presupposition. *Lingua* doi:10.1016/j.lingua.2006.05.006.
- Stone, Matthew. 1994. The reference argument of epistemic *must*. *Proceedings of the International Workshop on Computational Semantics (IWCS)* 1. 181-190. <http://www.cs.rutgers.edu/~mdstone/pubs/iwcs94.pdf>.
- Swanson, Eric. 2005. Something *Might* might mean. Ms, MIT.
- Swanson, Eric. 2008. Modality in language. *Philosophy Compass* 3(6). 1193-1207. doi:10.1111/j.1747-9991.2008.00177.x.
- Veltman, Frank. 1985. *Logics for conditionals*: University of Amsterdam dissertation.
- Werner, Tom. 2006. Future and non-future modal sentences. *Natural Language Semantics* 14(3). 235-255. doi:10.1007/s11050-006-9001-8.
- Westmoreland, Robert R. 1995. Epistemic *must* as evidential. *Proceedings of the Amsterdam Colloquium* 10(3). 683-702.
- Westmoreland, Robert R. 1998. *Information and intonation in natural language modality*: Indiana University dissertation.
- Willett, Thomas. 1988. A cross-linguistic survey of the grammaticalization of evidentiality. *Studies in Language* 12(1). 51-97.