

A solution to Karttunen’s Problem¹

Matthew MANDELKERN — *Massachusetts Institute of Technology*

Abstract. There is a difference between the conditions in which one can felicitously assert a ‘must’-claim versus those in which one can use the corresponding non-modal claim. But it is difficult to pin down just what this difference amounts to. And it is even harder to account for this difference, since assertions of ‘Must φ ’ and assertions of φ alone seem to have the same basic goal: namely, coming to agreement that $\llbracket \varphi \rrbracket$ is true. In this paper I take on this puzzle, known as *Karttunen’s Problem*. I begin by arguing that a ‘must’-claim is felicitous only if there is a shared argument for its prejacent. I then argue that this generalization, which I call *Support*, can explain the more familiar generalization that ‘must’-claims are felicitous only if the speaker’s evidence for them is in some sense indirect. Finally, I sketch a pragmatic derivation of *Support*.

Keywords: epistemic modals, indirectness of ‘must’, Karttunen’s Problem, strength of ‘must’.

1. Introduction

Compare (1) and (2):

(1) It must be raining out.

(2) It is raining out.

Intuitively, an assertion of (1) and an assertion of (2) have the same basic aim: they are both proposals to accept that it is raining out. Once an assertion of (1) has been accepted, interlocutors are disposed to accept the content of (2): that it is raining out. Thus (1) seems to be *as strong as* (2). But it does not seem to be *stronger* than (2): it is very strange to assert (1) after (2) is already accepted, as witnessed by the oddness of (3):

(3) ??It’s raining; and moreover, it must be raining.

This suggests that assertions of (1) and (2) carry the same basic information. Yet the conditions under which they can be felicitously asserted differ in subtle ways. Suppose that Jane is in a windowless room, and sees her colleagues come in with wet umbrellas. Then she can assert either (1) or (2). But now suppose that Jane is looking out a window at the rain. She can still assert (2), but an assertion of (1)—‘It must be raining out’—would be decidedly odd.

¹I am grateful to audiences at MIT, the 2015 University of Chicago Workshop on Modality and Subjectivity, the New York Philosophy of Language Workshop, Arché, and Hampshire College; to reviewers for *Sinn und Bedeutung 21*; and to Justin Bledin, David Boylan, Agnes Callard, Nilanjan Das, Brendan de Kenessey, Janice Dowell, Daniel Drucker, Vera Flocke, Irene Heim, Matthias Jenny, Angelika Kratzer, Daniel Lassiter, Rose Lenehan, Sarah Murray, Dilip Ninan, Jacopo Romoli, Bernhard Salow, Ginger Schultheis, Brett Sherman, Alex Silk, Daniel Skibra, Eric Swanson, and Roger White for very helpful comments and discussion. Special thanks to Kai von Fintel, Justin Khoo, Robert Stalnaker, and Stephen Yablo.

Generally there exists a systematic difference between the conditions in which one can felicitously assert a ‘must’-claim with complement φ , versus the conditions in which one can felicitously assert φ alone.² This puzzle, known as *Karttunen’s Problem*,³ gets to the heart of a number of broad foundational questions concerning the meaning of epistemic modals and the structure of conversation, and will be the topic of this paper.

The argument of the paper comes in three parts. In §2, I get clear on the data: what exactly the difference in felicity conditions between sentences like (1) and (2) amounts to. The main claim in the literature, which I call *Indirectness*, is that a ‘must’-claim is felicitous only if the speaker’s evidence for its prejacent is *indirect*, whereas its bare prejacent can be asserted whether the speaker’s evidence is direct or indirect. I argue that, while *Indirectness* is correct, there is another, equally important, generalization which plays a key role in solving Karttunen’s Problem: namely, that a ‘must’-claim is felicitous only if the speaker ensures there is a salient argument in support of the claim’s prejacent. I call this constraint *Support*. In §3, I show that once we have *Support* clearly in sight, we can derive *Indirectness* through general pragmatic reasoning. In §4, I give a Gricean account of why *Support* arises in the first place. This solution to Karttunen’s Problem predicts that assertions of ‘Must φ ’ and φ have the same basic update effect, while explaining why the former, but not the latter, requires that the speaker share an argument for $\llbracket\varphi\rrbracket$, and that her evidence for $\llbracket\varphi\rrbracket$ be indirect.

2. The data

The main extant claim regarding the difference between an assertion of ‘Must φ ’ versus an assertion of φ alone is that it amounts to an *Indirectness* constraint:⁴

Indirectness: A claim of ‘Must φ ’ is felicitous only if the speaker’s evidence for $\llbracket\varphi\rrbracket$ is indirect; a claim of non-modal φ can be felicitous whether the speaker’s evidence for $\llbracket\varphi\rrbracket$ is direct or indirect.

Indirectness is motivated by considering pairs of sentences like (4a) and (4b):

- (4) a. It must be raining.
b. It’s raining.

First suppose that the speaker of (4a) is looking out at the rain. Then her assertion is distinctly odd; whereas if she said (4b), it would be unmarked. By contrast, if she has a piece of indirect evidence for rain—say, wet umbrellas—then either is felicitous. *Indirectness* is the most natural generalization to draw from data like these.

Indirectness partly characterizes the difference in felicity between ‘Must φ ’ and φ , but does it exhaust that difference? Most of the literature on Karttunen’s Problem has indeed focused exclusively on *Indirectness*. But a different, mostly neglected, thread has pointed to a further

²A ‘must’-claim is a claim containing an unembedded strong epistemic necessity modal.

³Following von Stechow and Gillies (2010), who credit Karttunen (1972) with bringing the issue to attention.

⁴Karttunen (1972), Veltman (1985), Kratzer (1991), von Stechow and Gillies (2010), Kratzer (2012), Matthewson (2015), Lassiter (2016), Giannakidou and Mari (2016), Sherman (2016).

contrast in felicity conditions between $\lceil \text{Must } \varphi \rceil$ and φ : in making a ‘must’-claim, the speaker must ensure that an *argument* for its prejacent is salient to all the interlocutors.⁵

Support: A claim of $\lceil \text{Must } \varphi \rceil$ is felicitous only if there is an argument for $\llbracket \varphi \rrbracket$ salient to all the interlocutors; a claim of non-modal φ can be felicitous whether or not there is a salient argument for $\llbracket \varphi \rrbracket$.

The data that motivate *Support* are less clearcut than those that motivate *Indirectness*. This is unsurprising: evaluating *Support* requires evaluating discourses as a whole, rather than single utterances, and it can be difficult to determine, in a given context, whether an argument has been made salient. In the remainder of this section, I will provide new data to argue that *Support* is indeed required to account for the difference in felicity conditions between $\lceil \text{Must } \varphi \rceil$ and φ .

Consider the following case:

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

Suppose the conversation ends here, and assume that Bernhard doesn’t know anything about Patch’s set-up at Mark’s house, or anything else which might help him figure out why Mark thinks that Patch was in the box of hay. There is something distinctly odd about this exchange. Intuitively, what Mark has said needs elaboration; either Mark should have proffered reasons to think that Patch was in the hay box, or Bernhard should have asked him for reasons, perhaps with, ‘Why do you say that?’ Here is a more felicitous version of (5); assume the same setup:

- (6) a. [Mark:] She’s great. I heard a snuffling from the box of hay on my way out—she must have gotten into the box.
- b. [Bernhard:] Cute!

Now suppose the conversation ends here. This exchange has none of the peculiarity of (5). Likewise, a non-modal variant of (5) is perfectly fine:

- (7) a. [Mark:] She’s great. She got into the box of hay this morning.
- b. [Bernhard:] Cute!

The infelicity of (5) thus seems to be due to the fact that a ‘must’-claim is made while no argument for its prejacent is given.

Cases like this provide our first piece of evidence for *Support*. To get another case of this kind on the table, consider (8), adapted from Murray (2014):

- (8) On her way to a meeting in a windowless building, Sarah sees Jim enter with a wet

⁵See especially Stone 1994.

umbrella. Sarah enters the meeting. Thomas, who didn't see the umbrella, asks 'What's the weather like?' Sarah responds:

- a. It must be raining out.
- b. It's raining out.
- c. It must be raining out; I just saw Jim come in with a soaking wet umbrella.
- d. It's raining out; I just saw Jim come in with a soaking wet umbrella.

Thomas replies: 'Oh, too bad. Ok, let's talk about the agenda for this meeting.'

If the conversation ends here, then (8a)—the variant with a 'must' and no argument—is odd, while the other variants are fine—again, just as *Support* predicts.

Comparing 'must' with other words that might at first glance seem to work in a similar way can help bring out the plausibility of *Support*. Consider (9), adapted from a television spy drama:

- (9)
- a. The suspect is fleeing south. We've sent agents ahead to Mattapan.
 - b. Why Mattapan?
 - (i) Apparently the Russians have a safe-house there.
 - (ii) The Russians must have a safe-house there.

If the conversation ends here, then (9b-ii) is peculiar in a way that (9b-i) isn't. 'Apparently', like 'must', is constrained by a form of *Indirectness*; but 'apparently', unlike 'must', is acceptable here without an argument. *Support* predicts precisely this contrast (assuming that no corollary governs 'apparently').

Finally, I note that the patterns reported here are robust across strong epistemic necessity modals, in English and in all the other languages I have checked.⁶

I thus conclude that we should adopt *Support* as part of our characterization of the difference in felicity conditions between a 'must'-claim and its bare prejacent. Before moving on, let me say more about what *Support* amounts to. First, what does an *argument* amount to? I will think of an argument for $[\varphi]$ in a particular context as a set of propositions which the speaker is commonly recognized to believe provides reason to believe $[\varphi]$ —either by deductively entailing its conclusion; by inductively supporting the conclusion; or by showing how the conclusion follows from what is already accepted.

Second, what does 'salience' amount to? I won't say much about this, but a few features are worth noting. First, an argument need not itself be commonly accepted (i.e. *common ground* (Stalnaker, 1970)). One can felicitously assert an argument conjoined with a 'must'-claim, even if the argument has not yet been (and never is) accepted by all the speakers (if Bernhard doesn't believe me that I heard a snuffling from the box of hay, this does not render (6) infelicitous). The sense in which an argument Γ must be salient is rather that it must be common ground that the speaker takes Γ to provide reason to believe the prejacent of her 'must'-claim, and that she is proposing to add Γ to the common ground. I will refer to an argument with this status as

⁶Bengali, French, German, Hindi, Japanese, Russian, Spanish, Swiss German, and Turkish.

‘salient’ or ‘shared’ or ‘publicly available’.

An important point about salience is that an argument can be salient without being made explicit. If Bernhard and Mark can both hear snuffling from the box of hay, then Mark can say ‘Patch must be in the hay box’, without any further argument. Here, the premise that merits Mark’s conclusion—that Mark can hear snuffling from the box—is sufficiently salient, rendering the ‘must’-claim acceptable. Similarly, arguments can sometimes be *accommodated*, provided it is clear enough from the context what the speaker has in mind.

Another noteworthy feature of salience is that the argument in question need not be salient *at the time of the assertion*; it can be provided shortly after the assertion, as in (10):⁷

- (10) a. [Mark:] Patch must have gotten into the box of hay.
b. [Bernhard:] Why do you say that?
c. [Mark:] I heard her snuffling around when I was leaving.

3. Explaining *Indirectness* via *Support*

Support, in addition to *Indirectness*, is thus necessary to characterize the difference in felicity conditions between a ‘must’-claim and its bare prejacent.⁸ How should we explain these differences? There are three possible strategies: account for *Indirectness* and *Support* separately; account for *Support* in terms of *Indirectness*, and give an independent account of *Indirectness*; or account for *Indirectness* in terms of *Support*, and give an independent account of *Support*. I will briefly arguing against the first two strategies and then pursue the third.

3.1. Against the first two strategies

Considerations of theoretical parsimony tell against the first strategy. What about the second? This strategy is *prima facie* attractive, since there are a number of extant attempts to give an independent account of *Indirectness*; it is natural to try to recruit them to explain *Support*. But there are two significant problems with this approach. The first is that I do not believe that extant accounts of *Indirectness* are satisfying. I will not make this case here, however, because a second, simpler point suffices to show that the second strategy is wrong-headed: there does not seem to be any way to reduce *Support* to *Indirectness*.

A natural first thought is that this reduction would go by way of a general pragmatic constraint that requires a speaker to share her evidence for a claim if that evidence is indirect. But there is no such pragmatic constraint, as we saw in cases above where non-modal claims were felicitous without shared evidence. A closely related thought is that there is a general pragmatic constraint which requires a speaker to share her evidence if she explicitly indicates the source of her evidence.⁹ But, again, as the example above with ‘apparently’ shows—and as cross-linguistic

⁷As for other constructions that require something to be made salient, like anaphora resolution.

⁸I will not take a stand on the further question of whether they are jointly sufficient.

⁹Thanks to Justin Bledin for discussion.

work on evidentials suggests (see Murray (2014))—there is no such constraint: one can use ‘apparently’ or evidential marking without sharing what your evidence is.

A natural second thought is that *Support* reduces to a requirement to assure your interlocutors that *Indirectness* is satisfied. But this approach is not plausible, for a few reasons. First, the cases given above that are felt to be infelicitous without an argument—Patch in her box, Sarah in her windowless office building, the Russian safe-house—there is simply no reason to worry that the speaker’s evidence might not be indirect. Second, it is not generally true that whenever a formulation is constrained by a form of *Indirectness*, the speaker must habitually share her evidence in order to reassure her interlocutors that it satisfies the constraint in question: again, we saw this in (9) with ‘apparently’, which is governed by an *Indirectness* constraint, but which doesn’t require a shared argument.¹⁰ Finally, from a more theoretical standpoint, it is hard to see why an *Indirectness* constraint would ever directly yield an obligation to share one’s evidence: we are fairly charitable in assuming that speakers are complying with felicity conditions. For instance, if *Indirectness* were encoded as a presupposition, then, on a standard approach to presuppositions, it will be required that it be common ground that the speaker’s evidence for the prejacent is indirect. But in general interlocutors are perfectly happy to *accommodate* presuppositions,¹¹ leaving it puzzling why *Indirectness* would lead to a requirement for a speaker to share her evidence.

3.2. Deriving *Support* from *Indirectness*

I thus do not see a promising way for the second strategy to go. By contrast, I will argue now that the third strategy—deriving *Indirectness* from *Support*, and then giving an independent explanation of the latter—provides a satisfying solution to Karttunen’s Problem. In brief, the idea is as follows. I argue that an assertion of ‘Must φ ’ is a bid to update the common ground with $\llbracket\varphi\rrbracket$. *Support* then says that it is a proposal to do so *on the basis of an argument* Γ . General principles forbidding redundant assertions entail that $\llbracket\varphi\rrbracket$ should not follow from Γ in a way that is mutually recognized to be obvious. Finally, speakers are generally obligated to give their best argument for $\llbracket\varphi\rrbracket$ if they’re giving an argument for $\llbracket\varphi\rrbracket$ at all. It follows that, in order for an assertion of ‘Must φ ’ to be felicitous, $\llbracket\varphi\rrbracket$ should not follow in a mutually obvious way from the best argument a speaker of ‘Must φ ’ can have for $\llbracket\varphi\rrbracket$; in other words, a form of *Indirectness*.

The first step in our derivation is the assumption that an assertion of ‘Must φ ’ is *pragmatically strong* in the sense that it is just as strong as an assertion of φ . Following Stalnaker (1978), I assume that an assertion of φ is a proposal to update the common ground with $\llbracket\varphi\rrbracket$. Then we can spell out *Pragmatic Strength* as saying that an assertion of ‘Must φ ’ is, *inter alia*, a proposal to update the common ground with $\llbracket\varphi\rrbracket$:

Pragmatic Strength: An assertion of ‘Must φ ’ is as strong as an assertion of φ ,

¹⁰And, again, Murray (2014) likewise observes that grammatical evidential markers for indirectness do not give rise to any obligation to share one’s evidence.

¹¹See e.g. Lewis (1979), Stalnaker (2002) and many others.

in the sense that once the common ground is updated with $\llbracket \text{Must } \varphi \rrbracket$, it is updated with $\llbracket \varphi \rrbracket$.

Pragmatic Strength says that conversants do not typically leave open the possibility of $\llbracket \neg \varphi \rrbracket$ after accepting $\lceil \text{Must } \varphi \rceil$. To see its plausibility, note the weirdness of (11b) and (11c) as responses to (11a):

- (11) a. The gardener must be the murderer.
b. I concur. Moreover, the gardener is the murderer.
c. I concur. Let's bring him and the butler in to see if we can pin down which of them is the murderer.

Pragmatic Strength provides the most natural explanation of the infelicity of (11b) and (11c).

The second step is to note that in general, when a speaker tries to get her interlocutors to accept $\llbracket \varphi \rrbracket$ on the basis of an argument, the argument must be *non-redundant* in a relevant sense. Compare the two variants in each of (12) and (13):

- (12) a. I put Patch in her box this morning, and no one has let her out. So she's in her box.
b. ??I see Patch in her box. So she's in her box.
- (13) a. Clinton has amassed a majority of pledged delegates and superdelegates. So a woman will clinch the Democratic nomination!
b. ??Clinton will clinch the Democratic nomination. So a woman will clinch the Democratic nomination!

Without further justification for their repetitiveness, there is something pedantic or redundant about (12b) and (13b). By contrast, (12a) and (13a) are fine. The difference seems to be that in (12a) and (13a), there is enough epistemic space left between the argument in the first sentence and its conclusion in the second that its conclusion is not felt to be redundant. This intuition can be regimented as a norm against redundant assertions, along the following lines:

Non-Redundancy: A proposal to update the common ground with $\llbracket \varphi \rrbracket$ on the basis of an argument Γ is infelicitous if $\llbracket \varphi \rrbracket$ follows from Γ in a way that is mutually recognized to be obvious.

Non-Redundancy nicely captures the contrast between (12a) and (12b). The first is acceptable, since having put Patch in her box in the morning, together with no one else having let her out, does not, in an intuitive sense, obviously entail that Patch is in the box. The second is not, since it does follow in a mutually obvious way from seeing Patch in her box that she is in her box. Likewise for (13).¹²

¹²Note that *Non-Redundancy* does not forbid post hoc *support* for an assertion with a redundant argument; it is perfectly fine to justify oneself, if challenged, with 'Because I saw it'. What *Non-Redundancy* forbids is making an initial bid to update the common ground with something on the basis of an argument from which it follows in

The last step in our derivation says that a speaker must give the best argument for $[\varphi]$ that she has, if she's giving an argument for $[\varphi]$ at all. To see the plausibility of this constraint, consider (14):

- (14) John was at the Red Sox game and knows on this basis who won. He also read about the game in the *Boston Globe*.
- a. [Max:] Who won the game?
 - b. [John:] ?? The Red Sox, according to the *Globe*.

If John intends (14b) to answer Max's question, then there is something strange about it; we expect John to give his strongest evidence for the claim that the Red Sox won. In general, speakers are required to share the best piece of evidence they have for a claim, if they are sharing evidence at all. This follows naturally from a broadly Gricean vantage point on conversational dynamics. In (14b), John is violating Grice's Maxim of Quantity by failing to 'make his contribution as informative as is required (for the current purposes of the exchange)' (Grice, 1989). More precisely, the general lesson of cases like this is a corollary of the Maxim of Quantity which I call *Strongest Evidence*:¹³

Strongest Evidence: When a speaker aims to update the common ground with $[\varphi]$ on the basis of an argument Γ , she is obligated to do so by providing the *strongest* argument—the best piece of evidence—which she has for that claim.

We can now put these pieces together to derive *Indirectness* from *Support*. *Support* says that an assertion of $\ulcorner \text{Must } \varphi \urcorner$ is felicitous only if there is a shared argument for $[\varphi]$. *Pragmatic Strength* says that an assertion of $\ulcorner \text{Must } \varphi \urcorner$ is a proposal to update the common ground with $[\varphi]$. I will make the plausible further assumption that an assertion of $\ulcorner \text{Must } \varphi \urcorner$ is thus a proposal to update the common ground with $[\varphi]$ *on the basis of* a shared argument for $[\varphi]$ (this is an assumption that will fall out of the derivation of *Support* below). According to *Non-Redundancy*, $[\varphi]$ must not follow from that argument in a mutually obvious way. According to *Strongest Evidence*, that argument must constitute the best evidence the speaker has for $[\varphi]$. It follows that in order for a speaker to be able to felicitously assert $\ulcorner \text{Must } \varphi \urcorner$, $[\varphi]$ cannot follow in a mutually obvious way from the speaker's best piece of evidence for $[\varphi]$. In other words, the speaker's best evidence for $[\varphi]$ must be indirect, in the sense of indirectness relevant to evaluating whether an argument is felt to be redundant.

Put differently: Suppose a speaker has direct evidence (in the sense relevant to judgments about redundancy) for $[\varphi]$. If she were to assert $\ulcorner \text{Must } \varphi \urcorner$, then due to *Strongest Evidence* and *Support*, she would have to give that evidence as an argument on the basis of which she is proposing her interlocutors accept $[\varphi]$; but then she would be bound to violate *Non-Redundancy*. So if she has direct evidence for $[\varphi]$, she cannot assert $\ulcorner \text{Must } \varphi \urcorner$.

In sum, in asserting $\ulcorner \text{Must } \varphi \urcorner$, the speaker has to ensure there is a shared argument which

a mutually obvious way.

¹³See Faller (2012) for more careful discussion of how this kind of reasoning would go. To spell out *Strongest Evidence* in more detail, we need to be able to access a scale of evidential strength, according to which, say, direct perceptual evidence counts as stronger than any kind of testimonial evidence.

represents her best evidence for $\llbracket \varphi \rrbracket$, and yet is not so strong that it makes the ‘must’-claim sound redundant. Thus $\llbracket \varphi \rrbracket$ can’t follow in a mutually obvious way from her best evidence for $\llbracket \varphi \rrbracket$. No parallel constraint follows for non-modal claims—since *Support* requires only that ‘must’-claims be supported by an argument—and thus *Support*, plus *Pragmatic Strength*, *Non-Redundancy*, and *Strongest Evidence*, entail a form of *Indirectness*.

3.3. Predictions

The present proposal derives *Indirectness* from general principles about redundant assertions, and thus makes a striking empirical prediction: namely, that S’s evidence Γ for $\llbracket \varphi \rrbracket$ counts as indirect in the sense relevant to *Indirectness* just in case an assertion of φ following sequential assertions of the elements of Γ does not strike us as redundant. In addition to providing a new explanation for *Indirectness*, this also provides a new characterization of the notion of indirectness involved, an answer which I will now argue provides a better characterization of the data than the natural alternative, according to which ‘must’ lexically encodes a requirement that the speaker’s evidence be indirect in a sense that lines up with intuitions about whether evidence is direct or indirect, and with categories which are encoded by grammatical evidentials (I’ll call this ‘an evidential approach’).

I will highlight a few points. First, the present approach predicts that ‘must’-claims based on reliable testimony like (15) will not be acceptable:

(15) ??The website says the movie is at 7:30. So the movie must at 7:30.

Our approach predicts this, since reliable testimony for $\llbracket \varphi \rrbracket$ is typically felt to be a redundant argument for $\llbracket \varphi \rrbracket$, as shown by examples like (16):¹⁴

- (16) a. What time is the movie?
b. ??The cinema website says that it’s at 7:30. So the movie’s at 7:30.

By contrast, this is surprising on an evidential approach, since testimony is, intuitively, indirect evidence (it is natural to say that you know that the movie is at 7:30, but that you know *indirectly*, via the website). An evidential approach thus must simply stipulate that testimony ‘counts as direct’ for the purposes of evaluating ‘must’.

The second prediction worth highlighting is that, on the present approach, what counts as redundant in a given context—and thus judgments about the felicity of ‘must’—depends on what counts as mutually obvious in that context. Thus, e.g., while (15) is infelicitous out of the blue, it may be felicitous in a context in which the inference from website listings to fact is not generally accepted, as in (17):

(17) Google says that the movie is at 7:30. Websites listing movie times are generally

¹⁴Why sequences like this are treated as redundant is, of course, an important question for theories of redundancy to address, but one I will not answer here.

extremely unreliable. Google is extremely reliable, though, so the movie is indeed at 7:30.

Given the felicity of (17), we predict that a ‘must’-claim will be felicitous here as well; and indeed, (18) is felicitous:

(18) Google says that the movie is at 7:30. Websites listing movie times are generally extremely unreliable. Google is extremely reliable, though, so the movie must indeed be at 7:30.

More generally, we rightly predict that judgments about the felicity of ‘must’-claims depend on what counts as mutually obvious in context. It is not as clear how an evidential approach would predict this, since it does not seem like what counts intuitively as direct versus indirect evidence varies from context to context: our evidence for the time of the movie is equally indirect, in an intuitive sense, in (18) as in (16).

Third, we can explain why ‘must’-claims that conclude a complicated argument are generally acceptable, even if the premises of the argument entail its conclusion. Examples of this kind, in particular those involving mathematical or logical claims, are the most puzzling examples for an evidential approach to ‘must’. ‘Must’ is often warranted in mathematical or logical contexts, like (19).¹⁵

(19) If the set of validities were decidable, then the halting problem would be decidable. The halting problem is not decidable. So the set of validities must be undecidable.

It is not clear what the evidential approach would predict about (19). It is not clear that there is an intuitive sense on which our evidence that the set of validities is undecidable is indirect. Perhaps the evidential approach would claim that evidence for mathematical claims is *always* indirect in the relevant sense. But in addition to being somewhat stipulative, this response runs into trouble when it comes to examples like (20):

(20) $24 + 24$ must be 48.

If all evidence from mathematics is indirect, then the evidential approach will wrongly predict that ‘must’ is warranted in (20) and sentences like it. In any case, it is not clear how the evidential approach can distinguish between cases like (19) versus (20). By contrast, our approach can. The conclusion of (19) does not follow in a way that is mutually obvious from the premises, whereas the conclusion of (20) does.

These points confirm the key claim of the present approach to *Indirectness*: what matters for

¹⁵Is the ‘must’ here epistemic? Some have argued that this ‘logical’ ‘must’ is not genuinely epistemic (e.g. Giannakidou and Mari (2016), Goodhue (2016)). Two things militate against this option. First, it is inelegant to multiply modal flavors further than we need to. Second, even if we say that the logical ‘must’ is not epistemic, we still need a theory of its distribution, since it is not always warranted, even when its complement is a logical consequence of the common ground (as examples like (20) show); simply saying that this ‘must’ is logical, not epistemic, thus does not yet explain its behavior.

determining whether a ‘must’ is warranted is not whether the speaker’s evidence for it is indirect in a sense which lines up with our intuitions about sources of evidence, or with categories encoded by grammatical evidentials, but rather whether the speaker’s evidence for it makes the prejacent *mutually obvious*.

I conclude by discussing a different prediction of the present account. I have proposed that *Indirectness* arises due to conversational norms. It is a hallmark of pragmatic phenomena like this that they can be cancelled, since the underlying conversational norms are generally defeasible. We thus predict that *Indirectness* will be cancelled when one of the underlying norms is not in play. This prediction, again, is borne out, in particular in contexts in which *Strongest Evidence* is not in play because it is overridden by considerations which prevent the speaker from sharing her strongest evidence for $[\phi]$. For instance, suppose that Mary is at Tom’s party. She goes out to the street to smoke, where she runs into Ben. She knows Ben wasn’t invited to the party, and doesn’t want him to know that she was invited. Ben can hear music coming from Tom’s place, and asks Mary what’s going on at Tom’s. Mary wants to communicate that he’s having a party, but she doesn’t want to share her strongest evidence for this—and doesn’t seem to be under any obligation to do so, since she is trying not to hurt Ben’s feelings. In this context, she can felicitously assert (21):

(21) Given the music, it must be some kind of party.

(21) may be misleading, but it is perfectly felicitous, despite the fact that Mary’s evidence is direct. The prediction of our pragmatic account is thus borne out: *Indirectness* can be violated when one of the underlying pragmatic norms can itself be appropriately ignored.¹⁶

4. Support

Support plus independently motivated pragmatic principles thus provide a satisfying explanation of *Indirectness*. I turn now to the question of how to account for *Support*. I briefly criticize extant proposals before giving my own account.

4.1. Extant proposals

Support says that ‘must’ requires that an argument for its prejacent be made salient. A natural first thought about how to account for *Support* is to treat ‘must’ as containing something like an implicit indexical which refers to an argument: ‘must’ means roughly ‘it follows from *this argument* that...’, where the implicit ‘this’ requires a salient referent. Stone (1994) suggests an account along just these lines: on his approach, ‘must’ has a lexical argument place which must be saturated by an argument made salient by context. In other words, ‘must’ denotes a two place operator, taking an argument and a proposition p , which says that the argument provides decisive reason to believe p .

¹⁶It is less clear to me whether similar cases can be constructed in which it is *Non-Redundancy* which is suspended, since *Non-Redundancy* already has an element of context-sensitivity built in (since what counts as ‘mutually obvious’—in the objectionable sense relevant here—is itself context-dependent). But if we can find contexts in which it is suspended, then we predict that in those contexts as well, *Indirectness* will be suspended.

But a solution along these lines, natural though it is, does not work. The issue stems from the fact that no parallel to *Support* shows up for unembedded epistemic possibility modals. Consider (22):

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

Suppose the conversation ends here. As *Support* predicts, (22c) is infelicitous as it stands, without an argument. By contrast (22a)—like the non-modal variant in (22b)—is perfectly fine here. This suggests that 'might' is not subject to a *Support*-like constraint.

If we took Stone's approach, however, then, provided we assume that 'must' and 'might' are duals, we would predict that 'might' has an anaphoric requirement for an argument, just as 'must' does: if 'might' means 'not must not', then the argument requirement of 'must' will project through negation, and thus 'might' will require a salient argument, too.

We could avoid this by giving up the assumption that 'must' and 'might' are duals, and that 'might' does not have a lexical argument place for an argument. But, crucially, going this way leads to a serious new puzzle. Assuming we treat 'cannot' as equivalent to 'not might', then we will predict that unembedded 'cannot' does not have an anaphoric requirement for an argument any more than unembedded 'might' does. But this is wrong: the same examples we used to motivate *Support* for 'must' above can be used to motivate it for unembedded 'cannot' (*modulo* obvious changes). Thus, for instance, consider (23):

- (23) Emma notices that her neighbor Phil hasn't taken in his mail in some time, and concludes that he is out of town. Another neighbor asks if Phil is around. Emma responds:
- a. No, he can't be.
 - b. No, he's not.
 - c. No, he can't be: no one has taken his mail in for a week.
 - d. No, he's not: no one has taken his mail in for a week.

The exchange ends here.

(23a) is marked as compared with the other variants in (23). As with the examples involving 'must', then, 'can't' seems to require that an argument be made salient, in this case not for its prejacent but for its negation.

Thus a lexical derivation of *Support* along the lines Stone suggests faces a dilemma: either treat 'must' and 'might/can' as duals, and wrongly predict that the latter have a *Support*-like requirement; or do not treat them as duals, and wrongly predict that 'cannot' lacks a *Support*-like requirement. This approach thus strikes me as a non-starter.

Similar criticisms extend to the account suggested in Swanson (2015), who builds on Kratzer (1981) in adopting a premise semantics for epistemic modals, with the added requirement that those premises be publicly available. This approach faces the same dilemma. If we treat ‘might/can’ as the duals of ‘must’, then this explanation will overgenerate: it will wrongly predict that they are likewise subject to *Support*, since they will likewise require a set of premises to be made public. Alternately, we could abandon duality, but then we cannot explain *Support* for ‘cannot’. The present objection can also be extended to a treatment of *Support* as a presupposition.¹⁷ Presuppositions project through negation; thus a presuppositional approach would either treat ‘might/can’ as duals of ‘must’, and thus wrongly predict that ‘might/can’ obey *Support*; or would abandon duality, and once again fail to predict *Support* for ‘cannot’.

4.2. *Support* as a manner implicature

We can avoid these problems by deriving *Support* as a pragmatic implicature along the following lines. The derivation depends in part on adopting a semantics for ‘must’ defended in Stalnaker (2014) and Mandelkern (2016).¹⁸ I will not try to motivate the semantics here in general terms, but here’s a sketch. The idea is that ‘must’ is a universal quantifier over the set of worlds compatible with what is common ground after the ‘must’-claim in question has been made and negotiated (either accepted or rejected; call this the *prospective common ground*). ‘Must φ ’ thus means, roughly, ‘We will commonly believe $[\varphi]$ after this claim is made and assessed’; ‘might’ is treated as the dual of ‘must’, and thus ‘Might φ ’ will mean ‘ $[\varphi]$ is compatible with what we commonly believe, after this claim is made and assessed’.¹⁹ The basic idea is the familiar one that ‘might’ and ‘must’ are used to coordinate on what structural properties the context set has. I adopt the present semantics for ‘must’ partly because I think it is plausible, and partly because it lends itself naturally to the present derivation. A similar derivation of *Support* may well be possible with a different underlying semantics for ‘must’; the crucial features of the semantics for present purposes are, first, that it is pragmatically strong; and, second, that it makes salient the question of the interlocutors’ collective doxastic relationship to its prejacent. Any semantics for ‘must’ with these two features will suffice for present purposes.

On the present approach, ‘Must φ ’ and φ alone have the same basic update effect: namely, adding $[\varphi]$ to the common ground. ‘Must φ ’ is thus in competition with a different assertion which has the same basic update effect but which is structurally simpler: namely, an assertion

¹⁷A suggestion due to Eric Swanson (p.c.).

¹⁸My approach and Stalnaker’s differ in some ways, but the differences do not matter for our purposes.

¹⁹This gloss is only rough, because these constructions embed differently. See Mandelkern (2016) for further discussion. The basic idea for embeddings is that the domain of quantification for embedded ‘must’ is determined by its *local context*. I will prescind from deciding whether to go in for a factive or a non-factive notion of common ground in spelling out this semantics.

of φ alone.²⁰ Because φ is structurally simpler,²¹ choosing an assertion like $\lceil \text{Must } \varphi \rceil$ instead requires some explanation. Because the two options have the same update effect, the interlocutors cannot reason that the speaker chose one of them because she didn't know the other, or knew the other was false (as in scalar reasoning). There is, however, an important difference between $\lceil \text{Must } \varphi \rceil$ versus φ alone: namely, that the former makes salient the question of the interlocutors' collective doxastic relation to $\llbracket \varphi \rrbracket$, while the latter does not: it is only about $\llbracket \varphi \rrbracket$. There are a variety of ways to make this intuition precise; it doesn't matter for our purposes how we choose between them.²² For our purposes, we can just treat the question raised by an assertion of any sentence ψ as the two cell partition $\{\psi, \bar{\psi}\}$. Then the question made salient by an assertion of $\llbracket \varphi \rrbracket$ is just the question whether $\llbracket \varphi \rrbracket$ is true; whereas the question made salient by an assertion of $\lceil \text{Must } \varphi \rceil$ is the question *whether we will come to commonly accept* $\llbracket \varphi \rrbracket$. Given that $\lceil \text{Must } \varphi \rceil$ has a simpler alternative with the same basic update effect (namely φ), if a speaker chooses to use this more complex expressions, then we will seek an explanation of this fact. Given that the chief difference between $\lceil \text{Must } \varphi \rceil$ and φ is in the question made salient by each, we will therefore reason that, in choosing the more complex option, she wishes to raise to salience the question of the group's collective doxastic relation to $\llbracket \varphi \rrbracket$, and thus the group's reasons for accepting $\llbracket \varphi \rrbracket$.

What kind of reason for accepting $\llbracket \varphi \rrbracket$ would be worth highlighting in this way? Whenever a speaker proposes to update the common ground with $\llbracket \varphi \rrbracket$, this provides *ceteris paribus* reason for her interlocutors to accept $\llbracket \varphi \rrbracket$: namely, the speaker's authority. A reason of this kind, therefore, is totally humdrum, and thus not worth highlighting. The use of an assertion like $\lceil \text{Must } \varphi \rceil$ rather than an assertion of φ alone thus can be justified only if the speaker wishes to highlight a substantial argument for $\llbracket \varphi \rrbracket$ —i.e., an argument over and above the speaker's authority—on the basis of which the speaker wishes her interlocutors to accept $\llbracket \varphi \rrbracket$. In other words (the interlocutors will reason), the speaker wishes them to accept $\llbracket \varphi \rrbracket$ on the basis of an argument that is commonly available to them. The speaker must, therefore, ensure that such an argument is salient—either by providing it, or being assured that her interlocutors can recover it from the common ground, possibly by accommodation.

In short: *Support* arises as a manner implicature, thanks to the fact that the speaker chose to propose to accept $\llbracket \varphi \rrbracket$ by way of an assertion which makes reference to the interlocutors' doxastic relation to $\llbracket \varphi \rrbracket$.

This approach has a number of attractions. First, it derives *Support* from a simple, independently motivated modal semantics, rather than lexical stipulation, and thus explains why *Support* arises for anything that has the meaning of 'must'. Second, it avoids the problem raised

²⁰See Degen et al. (2015) for a different manner-implicature based approach to explaining the behavior of 'must'. Like the present approach, that approach relies on the assumption that $\lceil \text{Must } \varphi \rceil$ is a costly alternative to φ whose use must be somehow explained. In contrast with the present approach, that approach attempts to derive *Indirectness* directly, rather than by way of *Support*, and without adverting to the different QUDs raised by modal vs. non-modal variants.

²¹I will not spell out assumptions about how we calculate which alternatives are relevant. It seems fairly plausible that however we do so, φ will count as a relevant alternative to $\lceil \text{Must } \varphi \rceil$ (and $\lceil \neg \varphi \rceil$ as a relevant alternative to $\lceil \text{Can't } \varphi \rceil$); this follows e.g. on the account given in Katzir (2007), according to which alternatives are calculated by the deletion or replacement of nodes at LF.

²²See Lewis (1988), Roberts (2012) and citations therein.

above for extant approaches. If we treat ‘might/can’ as duals of ‘must’, we will predict that the corollary of *Support* will be blocked for ‘might/can’, but will still be derived for ‘cannot’. First, the derivation of a *Support* constraint for ‘might/can’ will be blocked because a crucial step in our derivation of *Support* was that $\lceil \text{Must } \varphi \rceil$ has a structurally simpler alternative with the same basic update effect; this obviously does not hold for $\lceil \text{Might } \varphi \rceil$, which has a very different update effect than φ . And, second, assuming that ‘can’ and ‘might’ mean the same thing, and that ‘not’ scopes over ‘can’ in ‘cannot’, the derivation given above will extend immediately to unembedded ‘cannot’, predicting in particular that *Support* and *Indirectness* arise for the negation of the prejacent of ‘cannot’, just as they do for the prejacent of ‘must’.

Note that for the same reason that this derivation of *Support* is blocked for ‘might’, the corresponding derivation of *Support* will likewise be blocked for weak epistemic necessity modals like ‘ought’ and ‘should’, as well as probability modals, since assertions of $\lceil \text{ought/should/probably } \varphi \rceil$ and φ do not have the same basic update effect. These predictions again seem correct:

- (24) a. When do you want to meet?
 b. Let’s say Thursday;
 (i) I should be free then.
 (ii) I ought to be free then.
 (iii) I’ll probably be free then.
 (iv) I must be free then.

If the conversation ends here, responses (24b-i), (24b-ii), and (24b-iii) are all acceptable; by contrast, (24b-iv) is a strange way to end the conversation, and seems to require that some argument be given (‘...my secretary always leaves my Thursdays open’). The prediction of the present account, then—that modals which are not *Pragmatically Strong*, like ‘might/ought/should/probably’, do not carry a *Support* constraint—thus seems correct.²³

The present proposal makes a prediction which is worth highlighting: any construction which has the features which played a role in the present derivation is predicted to give rise to a *Support*-like constraint. That is, any construction which has φ as a structurally simpler relevant alternative; which has the same basic update effect as an assertion of φ ; and which highlights the speakers’ collective doxastic relation to $\llbracket \varphi \rrbracket$, is, *ceteris paribus*, predicted to give rise to *Support* (and thus also *Indirectness*). Further research should examine other expressions that share these three features to see if this prediction is borne out.²⁴

²³And the *Indirectness* inference for ‘might/ought/should/probably’ is easy to explain on pragmatic grounds without a detour through *Support*.

²⁴ $\lceil \text{It’s agreed that } \varphi \rceil$, $\lceil \text{Let’s agree that } \varphi \rceil$, $\lceil \text{We should believe } \varphi \rceil$, and $\lceil \text{It is clear that } \varphi \rceil$ all seem to have these properties, and do indeed appear to be governed by corollaries of *Support* and *Indirectness* (see Barker 2009 on the last of these). This suggests that, even if the pragmatic derivation I have given here is mistaken in some of its details, the explanation for both *Indirectness* and *Support* is pragmatic, and stems from the properties I have pointed to.

5. Conclusion

The argument of this paper has come in three parts. I began by arguing that, to fully characterize the differences in felicity conditions between an assertion of $\ulcorner \text{Must } \varphi \urcorner$ and an assertion of φ , we need not only *Indirectness* but also *Support*. Next, I argued that we can derive *Indirectness* from *Support* together with general pragmatic principles about assertions, and I argued that this derivation of *Indirectness* makes attractive predictions about when a ‘must’-claim is unacceptable. Finally, I made a proposal about how to derive *Support* as a manner implicature from a certain semantics for epistemic modals which treats them as quantifiers over the set of worlds compatible with the common ground.

The three parts of this argument are, to a degree, independent. If each of these moves is successful, however, then taken together, they constitute a solution to Karttunen’s Problem: characterizing and explaining the differences in felicity conditions between an assertion of $\ulcorner \text{Must } \varphi \urcorner$ and an assertion of φ . In short: because of its more complex form and the question it makes salient, $\ulcorner \text{Must } \varphi \urcorner$, unlike φ alone, requires that an argument be given for $\llbracket \varphi \rrbracket$; and from the requirement, in turn, we can conclude that the speaker’s evidence for $\llbracket \varphi \rrbracket$ is relevantly indirect.

In conclusion, I highlight a few broad upshots of my approach to Karttunen’s Problem. The first is about the relation between ‘must’ and evidentiality. I have argued that ‘must’ does not grammaticalize a certain constraint on the *type* of evidence, in any intuitive sense, which the speaker must have for its prejacent. Rather, the felt indirectness of ‘must’-claims is accounted for pragmatically, and it is accounted for not directly in terms of judgments about type of evidence but rather in terms of judgments about *redundancy*.

The second is about the meaning of ‘must’. My derivation of *Support* rests on a certain semantics for ‘must’. If the derivation is successful, then it provides an argument for adopting that semantics, and presents a challenge for advocates of different semantics for ‘must’: to show how those approaches can explain *Support*.

The third regards the theory of redundancy. I have argued that the norms that governs redundancy in assertions play a crucial role in explaining our interpretation of modal language. My sketch of what those norms are like, however, leaves unanswered substantial questions, in particular about what counts as a ‘mutually obvious’ inference—questions which I hope to explore in future work. A theory of redundancy—in essence, a theory of how our minds structure and access information—will play a central role in understanding the way that information is structured in discourse, and judgments about ‘must’-claims provide a rich source of data for this theory.

I close with an abstract point about the architecture of semantic and pragmatic theories. My proposal rests on the assumption that an assertion of $\ulcorner \text{Must } \varphi \urcorner$ and an assertion of φ have the same basic update effect, but different semantic values. Indeed, on the semantics I have sketched, $\ulcorner \text{Must } \varphi \urcorner$ and φ informationally entail one another—in the sense that a context updated with either one entails the other—but they do not semantically entail one another. The possibility of this divergence between update effect and meaning proved essential for simulta-

neously capturing the intuition that an update with $\lceil \text{Must } \varphi \rceil$ is pragmatically strong, and the intuition that it cannot always be asserted where φ alone can be. Not everyone thinks that we should distinguish between semantic content and pragmatic update effect at all; in particular, at a high level, certain threads in dynamic semantics aims to identify these. If the present approach to Karttunen's Problem is the right one, however, then distinguishing semantic content from pragmatic update effect in our theorizing about natural language turns out to be crucial.

References

- Barker, C. (2009). Clarity and the grammar of skepticism. *Mind and Language* 24(3), 253–73.
- Degen, J., J. T. Kao, G. Scontras, and N. D. Goodman (2015). A cost- and information-based account of epistemic *must*. 28th Annual CUNY Conference on Human Sentence Processing.
- Faller, M. (2012). Evidential scalar implicatures. *Linguistics and Philosophy* 35, 285–312.
- von Fintel, K. and A. Gillies (2010). Must...stay...strong! *Natural Language Semantics* 18(4), 351–383.
- Giannakidou, A. and A. Mari (2016). Epistemic future and epistemic MUST: nonveridicality, evidence, and partial knowledge. In *Mood, Aspect and Modality: What is a Linguistic Category?* University of Chicago.
- Goodhue, D. (2016). Epistemic *must* is not evidential, it's epistemic. In *Proceedings of the North East Linguistic Society (NELS)*, Volume 46.
- Grice, P. (1989). *Studies in the Way of Words*. Harvard.
- Karttunen, L. (1972). Possible and must. In J. Kimball (Ed.), *Syntax and Semantics*, Volume 1, pp. 1–20. New York: Academic Press.
- Katzip, R. (2007). Structurally-defined alternatives. *Linguistics and Philosophy* 30, 669–690.
- Kratzer, A. (1981). The notional category of modality. In H. Eikmeyer and H. Rieser (Eds.), *Words, Worlds, and Contexts: New Approaches in Word Semantics*. de Gruyter.
- Kratzer, A. (1991). Modality. In A. von Stechow and D. Wunderlich (Eds.), *Semantics: An International Handbook of Contemporary Research*, pp. 639–650. Berlin: de Gruyter.
- Kratzer, A. (2012). *Modals and Conditionals*. Oxford: Oxford University Press.
- Lassiter, D. (2016). *Must*, knowledge, and (in)directness. *Natural Language Semantics*.
- Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic* 8(1), 339–359.
- Lewis, D. (1988). Relevant implication. *Theoria* 54(3), 161–74.
- Mandelkern, M. (2016). How to do things with modals. Manuscript.
- Matthewson, L. (2015). Evidential restrictions on epistemic modals. In L. Alonso-Ovalle and P. Menendez-Benito (Eds.), *Epistemic Indefinites*. New York: Oxford University Press.
- Murray, S. (2014). Varieties of update. *Semantics and Pragmatics* 7(2), 1–53.
- Roberts, C. (1998/2012). Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics* 5, 1–69.
- Sherman, B. (2016). Open questions and epistemic necessity. Manuscript.
- Stalnaker, R. (1970). Pragmatics. *Synthese* 22, 272–289.
- Stalnaker, R. (1978). Assertion. In P. Cole (Ed.), *Syntax and semantics*, Volume 9, pp. 315–322. New York: Academic Press.
- Stalnaker, R. (2002). Common ground. *Linguistics and Philosophy* 25, 701–721.
- Stalnaker, R. (2014). *Context*. Oxford University Press.

- Stone, M. (1994). The reference argument of epistemic *must*. In *Proceedings of IWCS 1*, pp. 181–190.
- Swanson, E. (2015). The application of constraint semantics to the language of subjective uncertainty. *Journal of Philosophical Logic* 45(121), 121–146.
- Veltman, F. (1985). *Logics for Conditionals*. Ph. D. thesis, University of Amsterdam.