

What ‘Must’ Adds*

Matthew Mandelkern
All Souls College, Oxford

October 1, 2018

To appear, *Linguistics and Philosophy*

Abstract

There is a difference between the conditions in which one can felicitously use a ‘must’-claim like (1-a) and those in which one can use the corresponding claim without the ‘must’, as in (1-b):

- (1) a. It must be raining out.
b. It is raining out.

It is difficult to pin down just what this difference amounts to. And it is difficult to account for this difference, since assertions of $\lceil \text{Must } p \rceil$ and assertions of p alone seem to have the same basic goal: namely, communicating that p is true. In this paper I give a new account of the conversational role of ‘must’. I begin by arguing that a ‘must’-claim is felicitous only if there is a shared argument for the proposition it embeds. 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 propose a pragmatic derivation of *Support* as a manner implicature.

1 Introduction

Consider the following sentences:

*I thank two anonymous referees for this journal, and my editor, Paul Portner, for very helpful comments. Thanks to Joshua Knobe and Jonathan Phillips for invaluable help running and analyzing the experiment reported here, and to Justin Khoo and MIT for financial support for the experiment. This paper expands on Mandelkern 2017b, and I also thank my referees and audience at *Sinn und Bedeutung 21*; audiences at MIT, the University of Chicago, NYU, St. Andrews, and Hampshire College; and Justin Bledin, David Boylan, Agnes Callard, Nilanjan Das, Janice Dowell, Daniel Drucker, Kai von Fintel, Vera Flocke, Irene Heim, Matthias Jenny, Brendan de Kennesey, Justin Khoo, Angelika Kratzer, Daniel Lassiter, Rose Lenehan, Sarah Murray, Dilip Ninan, Jacopo Romoli, Bernhard Salow, Ginger Schultheis, Brett Sherman, Alex Silk, Daniel Skibra, Robert Stalnaker, Matthew Stone, Eric Swanson, Roger White, and Stephen Yablo for very helpful comments and discussion.

- (2) a. It must be raining out. von Fintel and Gillies 2010
b. It is raining out.

Intuitively, an assertion of (2-a) and an assertion of (2-b) have the same basic communicative aim: namely, to communicate that it is raining out. Once an assertion of (2-a) has been accepted, interlocutors are disposed to accept that it is raining out, as witnessed by the oddness of (3):

- (3) ??It must be raining; and moreover, it is raining.

Thus (2-a) seems to be as strong as (2-b). But it does not seem to be stronger than (2-b): it is very strange to assert (2-a) after (2-b) is already accepted, as witnessed by the oddness of (4):

- (4) ??It's raining; and moreover, it must be raining.

This suggests that assertions of (2-a) and (2-b) carry the same basic information. Yet the conditions in 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 (2-a) or (2-b). But now suppose that Jane is looking out a window at the rain. She can still assert (2-b), but an assertion of (2-a)—‘It must be raining out’—would be decidedly odd. Generally speaking, there is a systematic difference between the conditions in which one can felicitously assert a ‘must’-claim with complement *p*, versus the conditions in which one can felicitously assert *p* alone.¹ Accounting for these differences is a challenge known as *Karttunen’s Problem*;² the goal of this paper is to propose a solution to this problem.

The argument of the paper comes in three parts. I begin by trying to get clear on the data: what the difference in felicity conditions between sentences like (2-a) and (2-b) 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, suitably spelled out, correct, there is another, equally important, generalization which plays a key role in solving Karttunen’s Problem: namely, that a ‘must’-claim is significantly degraded unless the speaker ensures there is a salient argument in support of the claim’s prejacent. I call this constraint *Support*. Then I show that once we have *Support* clearly in sight, we can derive *Indirectness* through general pragmatic reasoning. This reasoning reduces our judgments about the indirectness of ‘must’ to judgments about when a sequence of assertions is redundant. Finally, I give a Gricean account of why *Support* arises in the first place.

¹A ‘*must*’-claim is a claim containing an unembedded and unrestricted strong epistemic necessity modal such as ‘must,’ ‘it can’t be’, etc., read epistemically. I use ‘must’ as an exemplar of such modals, but my claims here are about strong epistemic necessity modals in general. I use ‘*p*’ as a sentence variable; ‘*p*’ denotes the proposition expressed by *p* (suppressing implicit relativization to contexts for readability).

²Following von Fintel and Gillies (2010), who credit Karttunen (1972) with bringing the issue to attention.

2 Data

I begin by trying to get clear on the difference in felicity conditions between assertions of \lceil Must p \rceil and p .

2.1 *Indirectness*

The main claim in the literature is that the difference amounts to an *indirectness constraint*:³

Indirectness: Asserting \lceil Must p \rceil is felicitous only if the speaker’s evidence for p is indirect.

By contrast, of course, a non-modal claim can be felicitous whether the speaker’s evidence is direct or indirect. *Indirectness* is motivated with cases like (5) and (6):

- (5) [Watching the rain:]
a. ??It must be raining.
b. It’s raining.
- (6) [Seeing your colleagues enter with wet umbrellas:]
a. It must be raining.
b. It’s raining.

(5-a) is distinctly weird as compared with (5-b), (6-a), and (6-b). *Indirectness* is the most natural generalization to draw from data like these. *Indirectness* has been well-motivated in the literature, and so I will assume without further discussion that *some* precisification of the principle is correct. More, of course, needs to be said about what counts as “indirect” here, as well as about the source of *Indirectness*, both questions which I return to below (see §3.4 in particular). For now, though, note that I do not assume that the concept of indirectness which plays a role in *Indirectness* neatly matches pre-theoretic intuitions about whether evidence is direct or not.⁴ For instance, reliable testimony is (according to some intuitions) *indirect* evidence; but, as von Stechow and Gillies (2010) observe, we must treat it as “direct” when it comes to evaluating *Indirectness*, in order to predict the contrast between (7-a) and (7-b):

³The insight goes back to Karttunen 1972. See von Stechow and Gillies 2010; also Veltman 1985; Kratzer 2012a; Matthewson 2015; Ozturk and Papafragou 2015; Lassiter 2016; Sherman 2018. Giannakidou and Mari (2016); Goodhue (2017) both argue that this empirical generalization is better explained in terms of partial knowledge, rather than indirectness; but they both still, as far as I can tell, endorse the generalization when understood in a sufficiently broad way, as intended here. As I emphasize in a moment, I am leaving the notion of indirectness intentionally broad and vague here; in what follows, I will propose one particular way of cashing it out, a way which, of course, not all these writers would necessarily endorse.

⁴Or categories which are encoded as grammatical markers of evidentiality in some languages; see e.g. Willett 1988; Faller 2002; Aikhenvald 2004.

- (7) [Tom tells Susie that it is raining. Susie says to Mark:]
- a. ??It must be raining.
 - b. It's raining.

2.2 *Support*

Most of the literature on Karttunen's Problem has focused exclusively on *Indirectness*. But a different thread in the literature has pointed to a further contrast in felicity conditions: in making a 'must'-claim, the speaker must ensure that an argument for its prejacent is salient to all the interlocutors. I call this claim *Support*:

Support: 'Must p ' is degraded unless there is an argument for p salient to the interlocutors and endorsed by the speaker.

By contrast, 'Must p ' can be felicitous when an argument for p is given. And non-modal claims, or modal claims with weaker epistemic modals ('might', 'should', etc.) can be felicitous with or without an argument.

The key insight in *Support* is due to Stone (1994). The claim, however, has not received much discussion in the subsequent literature.⁵ And indeed, the data that motivate *Support* are somewhat less clearcut than those that motivate *Indirectness*. This is not surprising: it can be difficult to determine, in a given context, whether an argument has been made salient. This requires evaluating discourses as a whole, and it requires considering whether an argument has been introduced either by linguistic means or through (possibly tacit) shared background knowledge or accommodation. In the remainder of this section, I will adduce a variety of data to argue that *Support* is indeed required to account for the difference in felicity conditions between 'Must p ' and p , as well as the difference in felicity conditions between 'Must p ' and modal claims of other strengths.

I begin by considering a variety of cases which bring out the need for *Support*. Consider first (8):

- (8) [Patch the rabbit sometimes gets into the cardboard box where her hay is stored. On his way out the door, Mark hears a snuffling from the box and thinks to himself, 'Patch must be in the hay box.' When he gets to school, Bernhard asks him how Patch is doing.]
- 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 in general anything which might help him figure out why Mark thinks that Patch was in the box of hay. There is something distinctly odd about this

⁵It is taken up briefly in Murray 2014, Silk 2016, Swanson 2015, and Lassiter 2016.

exchange. Intuitively, what Mark has said needs more 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 by saying, ‘Why do you say that?’ Here is a more felicitous version of Mark’s assertion in (8):

- (9) She’s great. I heard a snuffling from the box of hay on my way out—she must have gotten into the box.

Now suppose the conversation ends here. This exchange has none of the peculiarity of (8).

What does this show? First, note that a non-modal variant of (8) is perfectly fine:

- (10) a. [As in (8), except Mark says:] She’s great. She got into the box of hay this morning.
b. [Bernhard:] Cute!

The non-modal variant on (9) is likewise fine. The infelicity of (8) thus seems to be due to the use of ‘must’. But note that Mark’s evidence regarding Patch’s whereabouts is exactly the same in (8) as in (9). Moreover, *Indirectness* is satisfied—Mark’s evidence about Patch’s whereabouts is indirect—so *Indirectness* cannot explain the contrast between them. By contrast, *Support* predicts precisely the pattern we observe: namely, that (8) will be relatively degraded—since no argument for the prejacent is made salient—while (9) and (10) will be more acceptable.

Cases like this provide our first piece of evidence for *Support*. Before giving further cases, I will say a bit more about what *Support* amounts to. First, what does an *argument* amount to in this context? I will treat an argument for p as a set of propositions which the speaker is commonly recognized to believe provides reason to believe p —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 common ground, i.e. need not itself be commonly accepted in the conversation, to count as salient.⁶ 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 (9) 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 ‘salient’, ‘shared’, or ‘publicly available’; I have in mind this somewhat technical sense throughout.

An important point about salience is that an argument can be salient without being made

⁶See e.g. Stalnaker 1970, 2002, 2014.

explicit, as in (11):

- (11) [Bernhard and Mark are in the bunny’s room, and can both hear snuffling from the box of hay. Mark:] Patch must be in the hay box.

Here, the premise that merits Mark’s conclusion—that Mark can hear snuffling from the box—is salient, and the ‘must’-claim is acceptable.

Another noteworthy feature of the notion of salience in question is that the argument need not be salient *at the time of the assertion*; it can be provided shortly after the assertion, as in (12):

- (12) a. [As in (8), but Mark says] 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.

A final important fact about the salient argument in question is that it must be endorsed by the speaker of the ‘must’-claim. To see this, compare the following:

- (13) a. Patch must have gotten into the box of hay; John told me that he heard her snuffling around in there.
b. ??Patch must have gotten into the box of hay; John told me that he heard her snuffling around in there, though John often lies about things like that.

An argument which is salient but not endorsed, as in (13-b), clearly does not satisfy *Support*.

2.3 More cases

We find further confirmation of *Support* when we turn our attention to a broader range of parallel cases. Consider (14), adapted from Murray 2014:

- (14) [Sarah works in a windowless building. On her way to a meeting, she sees her coworker Jim enter the building, carrying a wet umbrella. Sarah concludes from this that it’s raining out. Sarah enters the meeting. Her colleague Thomas, who didn’t see Jim carrying a wet 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.]

As *Support* predicts, (14-a)—the variant with ‘must’, but without an argument—is odd, while

the other variants ('must' with an argument, and non-modal claims with or without an argument) are felicitous. We find a similar pattern in (15):

- (15) [Jane is in her first year of college. She doesn't have a clear sense of how she is doing in school. She meets with her professors, who tell her she is doing well; she thus concludes that she is doing okay. She goes in to meet with her adviser to talk about course registration. Her adviser doesn't know about the conversations she's had with her professors. Her adviser asks: 'So, how are you doing in your classes?' Jane responds as follows:]
- a. ??I must be doing okay!
 - b. I'm doing okay!
 - c. I must be doing okay: I've spoken to all my professors and they told me I'm doing fine.
 - d. I'm doing okay: I've spoken to all my professors and they told me I'm doing fine.

[Her adviser replies: 'Good, I'm happy to hear that. Ok, on to our business for today: let's discuss your registration for next term.']

Again, as *Support* predicts, (15-a) is marked, while the other variants are fine.

The next example illustrates the contrast of 'must'-claims without an argument not only to 'must'-claims with an argument and non-modal claims with or without an argument, but also to epistemic modal claims of other strengths (responses are labeled for ease of reference):

- (16) [Two friends, Scott and Mark, are discussing summer plans. Scott asks Mark: 'Do you think you'd be free to go fishing in a few months, say in the first week of September?' Mark responds:]
- a. [-arg, +must] ?? Yeah, I must be off work that Monday. Where would you want to go?
 - b. [+arg, +must] Yeah, that Monday is Labor Day, so I must be off work. Where would you want to go?
 - c. [-arg, -modal] Yeah, I'm off work that Monday. Where would you want to go?
 - d. [+arg, -modal] Yeah, that Monday is Labor Day, so I'm off work. Where would you want to go?
 - e. [-arg, +should] Yeah, I should be off work that Monday. Where would you want to go?
 - f. [+arg, +should] Yeah, that Monday is Labor Day, so I should be off work. Where would you want to go?
 - g. [-arg, +might] Yeah, I might be off work that Monday. Where would you want to go?
 - h. [+arg, +might] Yeah, that Monday is Labor Day, so I might be off work. Where

would you want to go?

- i. [-arg, +probably] Yeah, I'm probably off work that Monday. Where would you want to go?
- j. [+arg, +probably] Yeah, that Monday is Labor Day, so I'm probably off work. Where would you want to go?

Assuming that Mark doesn't remember that the first Monday in September is Labor Day, the 'must'-claim in (16-a) without further explanation is marked, whereas all the other variants—'must' with an argument, as well as non-modal, 'should', 'might', and 'probably'—are perfectly acceptable as is.

Two further examples illustrate an important point regarding *Support*: an argument can be shared even if it is not spelled out explicitly, but instead is quietly accommodated. (17) is taken from a radio show:

- (17) Mozart wrote the Stadler quintet for his friend Anton Stadler, who must have been a marvelous clarinetist.

The announcer does not give an explicit argument that Anton Stadler was a good clarinetist, but it is easy to recover from context. Second, suppose my phone rings; I can say:

- (18) This must be my brother; let me take this.

It is, again, easy for you to recover my reasons for saying this from the context (I must be expecting a call from my brother). (Some subtlety is required here. In general, when someone asserts 'Must p ', you can always thereby gather from that some reason to believe p , namely, that the speaker has proposed to update with p . So the contrasts we have seen suggest that the argument needed to satisfy *Support* must be more substantive than just *that the speaker believes p* . I return to this point in §4.2.)

Another piece of evidence for *Support* comes from naturalistic written examples. Lassiter (2016) discusses a range of examples taken from genealogical discussion boards, noting that 'Ancestry.com users frequently provide an explicit specification of the evidence used to arrive at a *must* conclusion'. Here's one example, taken from *The Plymouth Colony Archive Project*, slightly altered to improve flow across the variants below:

- (19) [+arg, +must] Goodman. . . is listed as one of those who received land in 1623 (PCR 12: 4). However, he is not listed among those who were part of the cattle division of 1627—the year we are interested in here; he must have died before then.

A variant which omits the argument for the 'must'-claim, as in (20-a), is felt to be missing something, if nothing more is said on the subject. By contrast, a non-modal variant with or without that material sounds fine, as do variants with 'might' and 'likely':

- (20)
- a. [-arg, +must] ?? Goodman. . . is listed as one of those who received land in 1623 (PCR 12: 4). However, he must have died before 1627, the year we are interested in here.
 - b. [-arg, -modal] Goodman. . . is listed as one of those who received land in 1623 (PCR 12: 4). However, he died before 1627, the year we are interested in here.
 - c. [+arg, -modal] Goodman. . . is listed as one of those who received land in 1623 (PCR 12: 4). However, he is not listed among those who were part of the cattle division of 1627—the year we are interested in here; he died before then.
 - d. [-arg, +may] Goodman. . . is listed as one of those who received land in 1623 (PCR 12: 4). However, he may have died before 1627, the year we are interested in here.
 - e. [+arg, +may] Goodman. . . is listed as one of those who received land in 1623 (PCR 12: 4). However, he is not listed among those who were part of the cattle division of 1627—the year we are interested in here; he may have died before then.
 - f. [-arg, +likely] Goodman. . . is listed as one of those who received land in 1623 (PCR 12: 4). However, he likely died before 1627, the year we are interested in here.
 - g. [+arg, +likely] Goodman. . . is listed as one of those who received land in 1623 (PCR 12: 4). However, he is not listed among those who were part of the cattle division of 1627—the year we are interested in here; he likely died before then.

Comparing ‘must’ with other words that might at first glance seem to work in a similar way, like ‘apparently’, can help bring out the plausibility of *Support*. Consider (21), adapted from a television spy drama:

- (21)
- a. The suspect is fleeing south. We’ve sent agents ahead to Mattapan.
 - b. Why Mattapan?
 - (i) ??The Russians must have a safe-house there.
 - (ii) Apparently the Russians have a safe-house there.
 - (iii) The Russians have a safe house there.
 - (iv) The Russians might have a safe-house there.
 - (v) The Russians probably have a safe-house there.

If the conversation ends here, then (21-b-i) is peculiar in a way that the ‘apparently’, non-modal, ‘might’, and ‘probably’ variants are not. ‘Apparently’ is of particular interest here, since ‘apparently’, like ‘must’, is constrained by a form of *Indirectness*; but ‘apparently’, unlike ‘must’, is acceptable here without an argument.⁷ This helps illustrate the separability of

⁷I assume that the indirectness constraint for ‘apparently’ is somehow lexically encoded. An anonymous referee for this journal rightly points out that this might provide support for a divided strategy, on which *Support* and

Indirectness from *Support*, a point I will return to.

Finally, I note that informal polling suggests the contrasts reported here seem to be robust across strong epistemic necessity modals in English (see §4.1 for further discussion of ‘can’t’) and other languages.⁸

2.4 Experimental results

The intuitions reported in the last section provide evidence for *Support*. Informal polling suggests these intuitions are robust. In this section, I present the results of an experiment which provides further support for these judgments.

301 participants were recruited from Amazon Turk (<https://www.mturk.com/>). All participants were based in the United States. We excluded from the analysis anyone who was not a native English speaker and anyone who completed the study in 30 seconds or less, leaving us with 273 participants. The stimuli in the experiment comprised the scenario in (16) along with the eight possible responses there, which, again, vary as to whether there was an argument or not (the *argument condition*) and as to whether a modal was used, and if so, which of three options—‘must’, ‘might’, or ‘should’—was used (the *modal condition*). Each participant received the scenario, plus one of the eight possible responses, and then was asked:

Please tell us how natural you think Mark’s response was on the scale below.

The participant responded to this prompt by dragging a marker on a continuous scale from 0 to 100, with ‘0’ labelled ‘completely unnatural’ and ‘100’ labelled ‘completely natural’.

Each participant saw four scenarios total, one from each modal condition, chosen at random within that condition as to whether or not it included an argument. The order of the scenarios was randomized. The scenario was designed so that all four modal conditions are reasonable to use as a response to Scott’s question, and it was designed so that Mark’s follow-up (‘Where would you want to go?’) makes clear that he is moving on from his answer to the question, suggesting that further information regarding his work schedule will not be forthcoming.

Support predicts that participants will find the [-arg, +must] condition degraded as compared with the seven other conditions ([+arg, +must], as well as all other modal conditions with or without an argument). Participant responses, which are summarized in Figure 1, show precisely this effect: participants found the [-arg, +must] condition substantially degraded (around

Indirectness have independent sources (which may coincide, as in ‘must’, and may diverge, as for ‘apparently’). But there is evidence that the status of the indirectness constraint for words like ‘apparently’ is different from the status of the corresponding constraint for ‘must’. For instance, although reliable testimony does not count as indirect for ‘must’, it seems to for ‘apparently’. Suppose John tells us that it’s hailing out. Then ‘Apparently it’s hailing out’ is perfectly fine.

⁸Informants who confirm the predicted contrast in Bengali, French, German, Hindi, Japanese, Russian, Spanish, Swiss German, and Turkish; more careful research needs to be done to more thoroughly explore the cross-linguistic picture.

the mid-point on the scale of naturalness judgments), as compared with all the conditions (which were around the three-quarter point on the scale of naturalness judgments).




Figure 1: Mean ratings by condition. Error bars show 95% confidence intervals.

Analysis of participants’ naturalness ratings revealed a significant interaction between the modal condition (‘must’ vs. ‘should’ vs. ‘might’ vs. non-modal) and the argument condition ($F(3, 809) = 11.2, p < 0.001$). To investigate the cause of this interaction effect, we separately analyzed participants’ ratings within each modal condition for the [+/-arg] conditions. Within the [-arg] condition, as predicted, participants found Mark’s response with ‘must’ to be substantially degraded—around the mid-point on the scale of naturalness (Mean = 57.20, $SD = 28.86$)—by contrast to all the other modal conditions, which they judged around the three-quarters mark on the scale of naturalness (t 's $> -6.08, p$'s $< .001, d$'s > 0.74). Things were entirely otherwise in the [+arg] condition, however, where participants found Mark’s response to be quite natural (around the three-quarters mark) in all modal conditions, including the ‘must’ condition. Specifically, participants found ‘must’ to be significantly more natural when an argument was included (Mean = 74.55, $SD = 23.51$) as compared to ‘must’ without an argument ($t(257.9) = -5.43, p < .001, d = 0.66$). Participants found the other modal conditions with an argument to be roughly equally good, again around the three-quarters mark; partici-

pants did not significantly distinguish ‘must’ with an argument from ‘might’ with an argument ($t(267) = -0.51, p = .611, d = 0.06$), though there was some preference for the ‘should’ and non-modal variants over the ‘must’ and ‘might’ variants ($t's > -3.97, p's < .001, d's > 0.48$). In no condition other than ‘must’ was there a contrast between the [-arg] and [+arg] conditions that was similar to the observed contrast in the ‘must’ condition; in the other modal conditions there was some variation between the [+arg] and [-arg] condition, but this variation was small compared with the corresponding variation in judgments in the ‘must’ condition (a +16.99 change in mean naturalness ratings for [+arg, +must] as compared with [-arg, +must], versus a -2.54 change for ‘might’, a +6.25 change for ‘should’, and a +8.75 change for non-modal).

These experimental results confirm the prediction of *Support*. They show that—as the intuitions elicited throughout this section suggest—‘must’ is felt to be significantly degraded without an argument. By contrast, non-modal claims, ‘should’ claims, and ‘might’ claims are all felt to be relatively acceptable without an argument—around the three-quarters point on the scale of naturalness, as opposed to a mid-point judgment for ‘must’. And, as expected, the ‘must’ variant significantly improved when an argument was added, to around the three-quarters point on the scale of naturalness—roughly the same judgment as for non-modal, ‘might’, and ‘should’ claims, with or without an argument. Note that the results here are graded. Speakers do not reject ‘must’ statements out of hand when there is no argument, even though they find them to be substantially less good than ‘must’-claims with arguments, or other kinds of claims with or without arguments. One possible explanation of this is that it is the effect of noise in the experimental set-up. It may be, for instance, that, in responding to the stimuli, participants are assuming that Scott is able to accommodate an argument for the relevant claim: as we have seen, an argument can become salient to a speaker’s audience through implicit clues or background knowledge, something which is difficult to control for in the experimental set-up. But it also may be that *Support* is indeed graded and in some cases relatively weak. I will not try to decide between these options here, proceeding simply with the conclusion that *Support* is a real signal which must be accommodated, one way or another, in our theory of ‘must’.

3 Explaining *Indirectness* via *Support*

I thus conclude that *Indirectness* and *Support* are both necessary to characterize the difference in felicity conditions between a ‘must’-claim and non-modal claims, as well as weaker modal claims.⁹ We now turn to the question of how to explain these data. There are three strategies to consider. We could account for *Indirectness* and *Support* separately; we could account for *Support* in terms of *Indirectness*, and give an independent account of *Indirectness*; or we could account for *Indirectness* in terms of *Support*, and give an independent account of *Support*. In the rest of the paper I will pursue the third strategy. In this section I will briefly argue against

⁹I will not try to settle the important further question of whether they are jointly sufficient.

the first two strategies, and then show how we can naturally account for *Indirectness* in terms of *Support*.

3.1 Negative remarks

Pro tanto considerations of theoretical parsimony tell against pursuing the first strategy (providing separate explanations for *Indirectness* and *Support*). What about the second strategy? 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 problems with this approach. Most significantly, there does not seem to be any way to reduce *Support* to *Indirectness*. A natural first thought is that we can explain *Support* in terms of *Indirectness* 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 in cases where the speaker's evidence was indirect but she did not share it. 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, in most of 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 no particular 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. We have seen this in (21) with 'apparently', which is governed by an *Indirectness* constraint, but which doesn't require a shared argument; likewise, weaker epistemic modals, like 'might', 'probably', and 'should', are governed by an *Indirectness* constraint (as von Stechow and Gillies (2010) emphasize), but, as we saw in (16) (and the corresponding experimental results), do not come with a *Support* constraint. Likewise, Murray (2014) observes that grammatical evidential markers for indirectness do not give rise to any obligation to share one's evidence. And from a more theoretical standpoint, it is hard to see why an *Indirectness* constraint would directly yield an obligation to share one's evidence. We are generally 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 often happy to accommodate presuppositions (Lewis, 1979; Stalnaker, 2002). There are caveats to this: for instance, interlocutors generally will not accommodate presuppositions which involve anaphoric content (Kripke, 2009), address a question under discussion, or contain noteworthy information (Stalnaker, 1974, 2002). But these caveats do not seem to apply in the present case. To derive *Support* from *Indirectness*, we would need an explanation of why interlocutors are not in general willing to accommodate that *Indirectness* is satisfied. The most important problem with the second strategy, then, is that there does not appear to be a promising way to derive

Support from Indirectness.

The second problem with the second strategy is that explaining *Indirectness* itself has turned out to be quite tricky: both main extant explanations have drawbacks which I will briefly survey here. I will not aim to argue decisively against those approaches, but rather to say enough to suggest that, even if we could reduce *Support* to *Indirectness*, there would remain open questions to address before we would have a satisfactory solution to Karttunen's Problem.

The first main approach to explaining *Indirectness* is pragmatic. On this approach, an assertion of $\lceil \text{Must } p \rceil$ proposes a weaker update than an assertion of p . By choosing the weaker update, the speaker signals that her evidence for p is too weak to merit an assertion of p . Thus we can conclude that the speaker's evidence for p is indirect, since direct evidence would have merited an assertion of p .¹⁰

This approach seems obviously correct for capturing the correlate to *Indirectness* for weaker modals like 'might', 'probably', and epistemic 'should' and 'ought'. Assertions involving one of these (unembedded) modals are uncontroversially weaker than the corresponding non-modal assertions, and thus this account provides a natural explanation of the corresponding indirectness inference. Nor, in those cases, does a requirement for an argument arise, as we have seen—so there is no question about how to extend this account of *Indirectness* to an account of such a requirement. But it is not clear that this approach extends equally well to the indirectness of 'must', because—as von Stechow and Gillies (2010, 2018) have forcefully argued—the assumption that $\lceil \text{Must } p \rceil$ is weaker than p is at odds with intuitions about a broad range of data involving 'must'. Those intuitions instead seem better captured by *Pragmatic Strength*:

Pragmatic Strength: An assertion of $\lceil \text{Must } p \rceil$ makes a proposal which is just as strong as an assertion of p .

I will follow Stalnaker 1978 in assuming that an assertion of p is a proposal to make the common ground entail p ; given that assumption, *Pragmatic Strength* says that an assertion of $\lceil \text{Must } p \rceil$ is as strong as an assertion of p in that both are proposals to make the common ground entail p . In other words, *Pragmatic Strength* says that conversants do not leave open the possibility that p is false after accepting $\lceil \text{Must } p \rceil$: to accept the latter is, *inter alia*, to accept the former. To see the motivation for *Pragmatic Strength*, consider the contrasts in (22):

- (22) [Detective A: 'The gardener must be the murderer.' Detective B, responding:]
- a. I concur. Let's arrest him.
 - b. ??I concur. Moreover, I've discovered that 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 actually is the murderer.

¹⁰The classic source of this approach is Karttunen 1972; Veltman 1985 and Kratzer 1977, 1981 give semantically weak semantics for 'must' which provide a basis for pragmatically weak accounts.

Pragmatic Strength explains the infelicity of (22-b) as a response to Detective A: if *Pragmatic Strength* is true, then p adds nothing over and above $\lceil \text{Must } p \rceil$, accounting for the infelicity of sequences of the form $\lceil \text{Must } p. \text{ Moreover } p \rceil$. And, likewise, *Pragmatic Strength* explains the infelicity of (22-c): once a group of interlocutors accepts $\lceil \text{Must } p \rceil$, then, if *Pragmatic Strength* is right, they should no longer leave open the possibility that p is false. By contrast, the infelicity of these sequences is hard to explain if *Pragmatic Strength* is false: if an update with $\lceil \text{Must } p \rceil$ leaves it open that p is false, then it seems like (22-b) and (22-c) should be perfectly coherent. This point can be brought out clearly by contrasting (22-b) and (22-c) with variants which use weaker modals which uncontroversially do not satisfy *Pragmatic Strength*. Suppose that, instead of saying ‘must’, Detective A said ‘The gardener is [probably/most likely/almost certainly] the murderer.’ Then the responses in (22-b) and (22-c) become perfectly felicitous. This is, again, exactly what we would expect if *Pragmatic Strength* holds for ‘must’, since it uncontroversially fails to hold for ‘probably’, ‘most likely’, ‘almost certainly’, and so on; and the contrast would be puzzling if ‘must’ were pragmatically weak.¹¹

Pragmatic Strength can also be motivated, following von Stechow and Gillies 2010, based on the infelicity of conjunctions of $\lceil \text{Must } p \rceil$ with conjuncts which propose to leave open the possibility that p is false, as in (23):

- (23)
- a. #The gardener must be the murderer, but there’s some chance that he isn’t.
 - b. #The gardener must be the murderer, but he might not be.
 - c. #The gardener must be the murderer, but let’s leave open the possibility that he’s not.
 - d. #The gardener must be the murderer, but the butler is also a possibility.
 - e. #The gardener must be the murderer, but it’s possible he’s not.
 - f. #The gardener must be the murderer, but I don’t know whether he is.
 - g. #The gardener must be the murderer, but I don’t know that he is.

The infelicity of these variants is, again, explained straightforwardly by *Pragmatic Strength*. If updating with $\lceil \text{Must } p \rceil$ entails an update with p , then it will be incoherent to propose an update with $\lceil \text{Must } p \rceil$ together with a proposal to leave open that p is false, as in (23-a)–(23-e). And, likewise, assuming that one should only propose an update with p if one knows p (as evidenced by Moore’s paradox (Moore, 1942); see Williamson 2000 for recent discussion), *Pragmatic Strength* accounts for the infelicity of (23-f)–(23-g). By contrast, it is hard to see how we would explain the infelicity of the variants in (23) if ‘must’ is not pragmatically strong.¹²

¹¹Indeed, it seems like for any $n < 1$, there is a clear contrast in this respect between $\lceil \text{There is an } n\% \text{ chance that } p \rceil$ versus $\lceil \text{Must } p \rceil$. This seems like a particular problem for degree semantics (like Lassiter (2016)’s) which predicts these to be Strawson-equivalent when n is sufficiently high.

¹²Lassiter (2016) adduces some naturalistic sequences of the form $\lceil \text{Must } p; \text{ possible not } p \rceil$, and nearby variations, which he argues are felicitous, contrary to my claim here. I find the data unconvincing, though I do not have space to discuss them in detail here. I believe that, in each case, the intended parse is *not* of the form $\lceil \text{Must } p; \text{ possibly not } p \rceil$. The principle argument for this is that, in every case, ‘must’ can be eliminated without sacrificing

All these considerations thus speak in favor of *Pragmatic Strength*, and at least suggest to me that *Pragmatic Strength* should be the default hypothesis. The present approach to *Indirectness*, which rejects *Pragmatic Strength*, incurs an as-yet unpaid debt: accounting for the striking data in (22)–(23). If *Pragmatic Strength* is correct, however, then of course the present derivation of *Indirectness* does not work.¹³

The main alternative explanation of *Indirectness*, due to von Stechow and Gillies (2010),¹⁴ accepts *Pragmatic Strength*, and posits that epistemic modals lexically encode that the speaker’s evidence is indirect as a semantic presupposition. This approach is motivated by, and avoids, the objection just given to the pragmatic approach. But this approach has its own drawbacks. The most serious of these, pointed out by Ippolito (2017), is that this approach does not make the right predictions about the projection of *Indirectness* from sentences which embed epistemic modals. I will not rehearse this point, which is somewhat complicated but strikes me as quite serious, here.¹⁵ My present aim, again, is not to close the door on either of the extant approaches to *Indirectness*, but rather to make the points, first, that both leave *Support* unexplained; and, second, that both have (perhaps serious) drawbacks; which together suggest that an alternative approach is worth exploring.

3.2 From *Support* to *Indirectness*

In the rest of this section, I explore just such a strategy, arguing that we can derive *Indirectness* from *Support*. In brief, the idea is as follows. First, *Pragmatic Strength* says that an assertion

coherence. If the intended reading were Lassiter’s, the results without ‘must’ would have the form ‘ $\Box p$ and possibly not $\Box p$ ’, which are uncontroversially incoherent (see Wittgenstein 1953; Groenendijk et al. 1996; Gillies 2001; Yalcin 2007). In any case, the comparison without ‘must’ suggests that these cases do not evidence a contrast in strength between ‘must’-claims and non-modal claims. Lassiter also reports an experiment in which participants are presented with a raffle scenario, and then asked to evaluate ‘Bill did not win the raffle’, ‘Bill must not have won the raffle’, ‘It is possible that Bill won the raffle’, and so on. Participants agreed with these three claims 69%, 58%, and 92% of the time, respectively. As Lassiter notes, these results are consistent with ‘ $\Box p$ ’ entailing p (and thus with *Pragmatic Strength*), since participants were just as inclined to accept p as they were to accept ‘ $\Box p$ ’.

¹³*Pragmatic Strength* is distinct from the strictly stronger principle that ‘ $\Box p$ ’ entails that contextually relevant agents *know* p (which Lassiter (2016); Giannakidou and Mari (2016) argue against), and from the strictly stronger principle that ‘ $\Box p$ ’ entails p (which Lassiter leaves open). Giannakidou and Mari (2016, 2017) defend an *obligatorily* weak view, on which ‘ $\Box p$ ’ is not only *compatible* with lack of knowledge of p , but in fact *entails* that the speakers do not have “complete knowledge” of p . This, however, strike me as untenable. If this were correct, then sequences of the form ‘ $\Box p$, and I know p for sure’, and nearby variants, should be infelicitous, but they are not, as witnessed by the felicity of sequences like (24):

- (24) John must be at home; I [am completely sure/am certain/know for sure] he’s there, because he’s bedridden and cannot leave on his own.

Goodhue (2017) similarly argues that ‘ $\Box p$ ’ is felicitous only if p is not known, which likewise seems to predict that (24) will be infelicitous, contrary to fact. (In places Goodhue seems sympathetic to the idea that this constraint is defeasible; the present point of course does not tell against that view.)

¹⁴And since (at least partly) endorsed and elaborated in Kratzer 2012a, Matthewson 2015, Lassiter 2016.

¹⁵See also Sherman 2018; Matthewson 2015 for further discussion of criticisms of the lexical approach.

of $\ulcorner \text{Must } p \urcorner$ is a bid to update the common ground with p . *Support* says that it is a proposal to do so on the basis of an argument Γ . General principles concerning redundancy entail that p 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 p if they're giving an argument for p at all. It follows that, in order for an assertion of $\ulcorner \text{Must } p \urcorner$ to be felicitous, p should not follow in a mutually obvious way from the best argument a speaker of $\ulcorner \text{Must } p \urcorner$ has for p . In short, an assertion of $\ulcorner \text{Must } p \urcorner$ is a proposal to accept p on the basis of a shared argument, and this is the kind of proposal one should make only when that argument leaves some epistemic space between its premises and conclusion.

The first step in our derivation is the assumption, defended in §3.1 above, that an assertion of $\ulcorner \text{Must } p \urcorner$ is pragmatically strong, in the sense that it is, *inter alia*, a proposal to update the common ground with p .

The second step is to note that in general, when a speaker gives an argument in support of p with the intention of getting her interlocutors to accept p on the basis of that argument, the argument in favor of p must be *non-redundant* in some sense. Compare the two variants in each of (25) and (26):

- (25) 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.
- (26) a. I've just read a CNN report that says Clinton has amassed a majority of pledged delegates and superdelegates. So Clinton will clinch the Democratic nomination!
 b. ??I've just read a CNN report that says Clinton will clinch the Democratic nomination. So Clinton will clinch the Democratic nomination!

(25-b) strikes me as objectionable; there is something pedantic or redundant about it. Likewise for (26-b), unless there is some salient doubt about the veracity of CNN. By contrast, (25-a) and (26-a) are fine. The difference seems to be that in (25-a) and (26-a), there is enough 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 p on the basis of an argument Γ is infelicitous if p follows from Γ in a way that is mutually recognized to be obvious.

In the next section I will further explore this principle. Note for the present, however, that *Non-Redundancy* nicely captures the contrast between (25-a) and (25-b), given our intuitive (albeit vague) notion of what it is for an inference to be mutually recognized to be obvious. The first is acceptable, since there is some epistemic space between the premises—having put Patch in

her box in the morning, together with no one else having let her out—and the conclusion—that Patch is in the box (epistemic space filled by background assumptions about Patch’s ability to get out on her own, etc.). 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. Similar considerations apply to the contrast between (26-a) and (26-b).

Note that *Non-Redundancy*, of course, does not forbid *post hoc* justification for an assertion with a redundant argument; it is perfectly fine to justify oneself, if challenged, with something like ‘Because I saw it’. What *Non-Redundancy* forbids is making an initial bid to update the common ground with *p* on the basis of an argument from which *p* follows in a mutually obvious way. Nor does *Non-Redundancy* forbid asserting *p* when *p* follows from your own private evidence in a mutually obvious way; see §3.3 for more discussion. Finally, note that *Non-Redundancy* is defeasible, as I discuss further in §5.3.

The last step in our derivation says that a speaker must give the best argument for *p* that she has, if she’s giving an argument for *p* at all. To see the plausibility of this constraint, consider (27):

- (27) [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 (27-b) 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 (27-b), 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 lesson of cases like this is a corollary of the Maxim of Quantity which I call *Strongest Evidence*:¹⁶

Strongest Evidence: When a speaker gives an argument for a claim, she must provide the best piece of evidence she can.

We can now put these pieces together to derive *Indirectness* from *Support*. *Support* says that an assertion of ‘Must *p*’ is felicitous only if there is a shared argument for *p*. *Pragmatic Strength* says that an assertion of ‘Must *p*’ is a proposal to update the common ground with *p*.

¹⁶See 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—on this point see also Faller 2001. *Strongest Evidence* is of course in tension with the aim of brevity, and ordinary conversational moves presumably aim to strike a balance between these.

It is natural to therefore assume that an assertion of $\lceil \text{Must } p \rceil$ is thus a proposal to update the common ground with p on the basis of a shared argument for p . According to *Non-Redundancy*, p 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 p . It follows that in order for a speaker to be able to felicitously assert $\lceil \text{Must } p \rceil$, p cannot follow in a mutually obvious way from the speaker's best evidence for p . In other words, the speaker's best evidence for p must be indirect, in the sense of indirectness relevant to evaluating whether an argument is felt to be redundant.

In sum: in asserting $\lceil \text{Must } p \rceil$, the speaker has to ensure there is a shared argument which represents her best evidence for p , and yet is not so strong that it makes the 'must'-claim sound redundant. Thus p can't follow in a mutually obvious way from her best evidence for p . No parallel constraint follows for non-modal claims—since *Support* requires only that 'must'-claims be supported by an argument—and thus *Indirectness* follows from *Support*, plus *Pragmatic Strength*, *Non-Redundancy*, and *Strongest Evidence*.

3.3 Non-Redundancy

Before turning to the question of how to predict *Support*, I will say more about the conversational architecture that underlies *Non-Redundancy* (in this subsection), as well as the predictions made by the present derivation of *Indirectness* (in the next subsection). More could be said about each of the assumptions I made in the last section, but *Non-Redundancy* is the one most in need of further explanation here, since a central upshot of my account is that the signal of indirectness associated with 'must' reduces to judgments about redundancy. Saying more about redundancy will help us evaluate the theoretical foundations and empirical ramifications of this claim.

Non-Redundancy forbids updating the common ground with p on the basis of an argument Γ when p follows from Γ in a way mutually recognized to be obvious. *Non-Redundancy* follows from two further claims. The first is that a proposal to update the common ground with p on the basis of an argument Γ is a proposal to update the common ground with Γ and then with p . The second is that one should not propose a series of updates if, should they all be accepted, the final update will be judged to be redundant. The first of these claims is fairly self-evident. The second requires further exploration. It has its foundation in the more general idea that there is something wrong with redundant assertions, along the following lines:

Common Ground Entailment: Don't assert p if p is entailed by the common ground.¹⁷

This norm is very natural if we think of conversations as cooperative enterprises whose main goal is information transfer. Asserting what is already common ground does not serve that goal.

¹⁷See Stalnaker 1974, 1978, 1973. Some version of this norm applies at the subsentential level as well; see e.g. Schlenker 2008, 2009, Mayr and Romoli 2016.

Common Ground Entailment, however, is not quite right.¹⁸ It overgenerates, for instance, in cases like (28) and (29):

- (28) a. [A:] If it's raining out, then Bob brought his umbrella. If Bob brought his umbrella, then he won't have noticed that we had the roof redone. And, it was raining out.
b. [B:] Ok.
c. [A:] So Bob won't have noticed that we had the roof redone.
- (29) a. [A:] If the set of validities were decidable, then the halting problem would be decidable.
b. [B:] Ok.
c. [A:] The halting problem is not decidable.
d. [B:] Right.
e. [A:] So the set of validities is not decidable.

Common Ground Entailment wrongly predicts that the last response in each of these sequences, which is entailed by what comes before, will be infelicitous. But in both cases, the final response is acceptable: it simply makes explicit the conclusion of a somewhat complex chain of reasoning. The problem, evidently, is that *entailment* is the wrong notion to play the role it does in *Common Ground Entailment*: sometimes assertions are not felt to be redundant even if they are contextually entailed, provided this is not mutually obvious.

A natural move at this point would be to say that an assertion is redundant if it is entailed by the common ground *in a way mutually recognized to be obvious* (in a derogatory sense: a sense which makes an assertion superfluous). In the present case, we say that (28-c) is *not* a mutually obvious entailment of (28-a), and thus we rightly predict its felicity; likewise, (29-e) is not a mutually obvious entailment of (29-a) and (29-b). What counts as a mutually obvious entailment will depend on the context, and will be a matter of ongoing negotiation, as I discuss further below.¹⁹

Merely changing 'entailed' to 'entailed in a mutually obvious way' does not yet, however, yield a satisfying principle. The modified principle will undergenerate in important cases. Consider (30):

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

¹⁸ A different place where *Common Ground Entailment* goes wrong is in terms of *side effects* like making a discourse referent available. See Barker and Taranto 2003; MacFarlane 2004; Barker 2009; Franke and de Jager 2007; Bronnikov 2008; Kripke 2009; Wolf and Cohen 2011; Crone 2015. I will, however, ignore side effects here for the sake of simplicity.

¹⁹ As Justin Khoo has pointed out to me, this will help explain why it can be rude to spell things out in detail: this shows that you are not treating something as mutually obvious which your interlocutor assumed was.

- c. Ok.
- d. ??So the movie's at 7:30.

(30-d) is, in most contexts, unacceptable—intuitively, because it is felt to be redundant. But (30-b) does not *entail* (30-d). To capture data like these, we need to use a more permissive relation than entailment: a relation which can capture the sense in which ‘the movie is at 7:30’ follows from ‘the cinema website says that the movie is at 7:30’.²⁰ Here I will simply use ‘follows from’ to capture this relation, with the understanding that we get a feel for the notion from examples like (30). If Γ entails p , then p follows from Γ ; but p also might follow from Γ if p is a reasonable default inference, or is highly probabilified by Γ . Much more needs to be said here in order to have a rigorous theory, of course, but this suffices for our purpose.

The most natural way to capture these two revisions to *Common Ground Entailment* is to model the common ground as a set of propositions which is closed under mutually obvious inference—if p is in the common ground and q follows from p in a mutually obvious way, then q is in the common ground—but which is not closed under logical entailment in general. Then we can state our redundancy principle simply as follows:²¹

Common Ground Settlement: Don’t assert p if p is in the common ground.

Common Ground Settlement, then, provides a reasonable basis for *Non-Redundancy*. The latter, in turn, follows from the former given the very natural assumption that one should not make an assertion which would violate a conversational norm were it accepted. This is required to account for cases like (25-b), where we cannot assume that ‘I see Patch in her box’ is actually common ground before ‘So she’s in her box’ is processed—only that it is intended to be.

There is, again, much more to be said about *Common Ground Settlement*—about when something follows from a set of premises, when it does so in a way mutually recognized to be obvious, and what underlying theory of mind is required to account for this. But as we will see now, even the first-pass account laid out here yields striking predictions about ‘must’.

3.4 Predictions

Our account reduces the sense of indirectness relevant to ‘must’-claims to the sense in which an inference can count as mutually obvious. Although much more could be done to spell out what mutually obvious inference amounts to, this does not mean that we have just reduced one

²⁰A different response to data like this would be to hold that in cases like this, we accommodate an enthymematic premise, so that this really is a contextual entailment. I do not see how to choose between this response and the present one; either formulation would do for our purposes.

²¹Note that if we model the common ground this way, then—if we follow standard assumptions in treating it as derivative from individual mental states (as in Stalnaker 2002)—the present considerations can be extended to argue in favor of likewise modeling individual attitudes as sets of propositions with these closure properties. The appropriate semantics for attitude reports would thus be a neighborhood semantics, rather than a standard modal semantics (Montague 1970, Scott 1970). See Kratzer 2012b, pp. 19-20 for related discussion.

obscure concept to another. For, while we don't have a complete theory of redundancy, we do have a rough operationalization of it: an inference from Γ to p is mutually obvious in a context just in case proposing to accept p on the basis of Γ in that context is felt to be redundant, otiose, pedantic, etc. This is obviously a bit vague, but the underlying notion is vague too, and, as I will now discuss, this characterization suffices to predict a variety of data which are either not predicted at all, or else are simply stipulated, on alternate accounts. In other words, having spelled out *Non-Redundancy* more carefully, we are in a position to flesh out a striking empirical prediction of our account: namely, that S's evidence Γ for p counts as indirect in the sense relevant to *Indirectness* just in case an assertion of p following sequential assertions of the elements of Γ does not strike us as redundant.

I will highlight a few points in this respect. First, as noted above, von Fintel and Gillies (2010) observe that reliable testimony generally counts as direct in the sense relevant to 'must':

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

This follows immediately on our approach. This is because reliable testimony for p is typically felt to make a subsequent assertion of p redundant, as shown by examples like (30) (a non-modal variant of (31)). *Why* sequences like this are treated as redundant is, of course, an important further question for theories of redundancy to address; but, without yet having an answer to this question, the present approach predicts this data point in a non-stipulative way.

The second prediction of our approach I will highlight is that what counts as redundant in a given context—and thus judgments about the felicity of 'must'—will be context-sensitive, since *Non-Redundancy* is context-sensitive. In particular, what counts as redundant in a context depends on what counts as mutually obvious in that context. Thus, e.g., while (31) 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 (32):

(32) 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 is indeed at 7:30.

If *Indirectness* bottoms out in judgments about redundancy, we thus predict that what counts as indirect will, likewise, be a context-sensitive matter. This prediction, again, is borne out. While (31) was strange out of the blue, it is much improved in (33), which, like (32), calls into question the reliability of websites listing movie times:

(33) 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.

A similar example, adapted from von Fintel and Gillies 2010, provides a nice further illus-

tration of this point. Suppose that we are at an epistemology conference, where the inference from appearances to reality have been constantly called into question. In such a context, the inference from ‘It appears to be raining’ to ‘It’s raining’ will not be treated as mutually obvious; sequences with the form of (34), while perhaps still somewhat odd, will be much improved as compared with ordinary contexts:

(34) It appears to be raining, so it is raining!

This leads us to predict that, in this context, our first example—using ‘must’ when we are looking out at the rain—will be much improved. And this is just what we find: (35) and (36) are much improved in this context, as compared with ordinary contexts.

(35) It appears to be raining, so it must indeed be raining!

(36) [Looking out at the rain] It must be raining!

Indirectness thus indeed seems to be context-sensitive in just the same way that *Non-Redundancy* is, as we predict.

These cases show a kind of context-sensitivity in which something like a skeptical scenario has been raised, making the inference from a given set of premises to a conclusion less secure than it normally is. This may make it look like the sense of indirectness relevant for ‘must’ patterns with knowledge ascriptions, something that Giannakidou and Mari (2016); Goodhue (2017) have indeed argued for. But as Goodhue points out, there are other forms of context-sensitivity relevant to ‘must’ that do not seem to have anything to do with knowledge. Here is one of Goodhue’s cases:

(37) [Billy is talking to her sister on the phone, and her sister has denied that it could be raining and has demanded repeatedly that Billy explain how she knows for sure that it is raining. At her wits’ end, Billy says:] If light enters your eyes in such a way that it looks like rain, then it is raining. Light is entering my eyes in such a way that it looks like rain. Therefore, it must be raining.

This is not a case of skeptical scenario raising—Billy’s sister’s obstinacy is not depriving Billy of knowledge, she is just being obstinate. Yet the ‘must’ in (37) seems felicitous. So the context-sensitivity of *Indirectness* does not, after all, pattern with the context-sensitivity of ‘knows’. In this case, however, it does again pattern with the context-sensitivity of redundancy: by refusing to accept ordinary inferential reasoning, Billy’s sister makes it clear that that reasoning is not *mutually* obvious (since it’s not obvious to her), and thus a ‘must’ here is licensed. Cases like this provide further support for the claim that the context-sensitivity of *Indirectness* patterns with the context-sensitivity of mutual obviousness (not with the context-sensitivity of knowledge ascriptions).

A third point I will highlight here is that the present account lets us explain why ‘must’-claims that conclude a complicated argument are generally acceptable, even if the premises of the argument entail its conclusion. ‘Must’ is often warranted in deductive contexts—mathematical or logical environments where explicit premises are given which entail the conclusion—like (38).²²

- (38) 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.

By contrast, in other deductive contexts, like (39), ‘must’ is not licensed:

- (39) I have three students in my logic class, and four other students in my semantics class, and those are my only classes. So I must have seven students total.

Our approach makes sense of this contrast in a straightforward way, since the conclusion of (38) does not follow in a way that is mutually obvious from the premises (even if it is entailed by them, and indeed by anything); whereas the conclusion of (39) does follow in a mutually obvious way from the premises.²³

A final set of predictions of the present account has to do with the structure of my derivation of *Indirectness*. I have proposed that *Indirectness* arises due to conversational norms. It is

²²Giannakidou and Mari (2016); Goodhue (2017) have argued that this ‘must’ is not genuinely epistemic, but instead is a “logical” or “deductive” ‘must’. It is hard to rule out an ambiguity thesis, but it is obviously inelegant to multiply modal flavors further than we need to. More seriously, even if we say that the ‘must’ in cases like these 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 relevant premises. That is, ‘must’ sounds fine in (38) and bad in (39). So, if we distinguish deductive from epistemic ‘must’, we end up doubling Karttunen’s Problem: we must then account for the distribution of ‘must’ in both epistemic and deductive contexts. Neither Giannakidou and Mari nor Goodhue offers a solution to this problem; e.g. if we simply say, following Goodhue, that the deductive ‘must’ ‘quantifies universally over the possibilities compatible with the overtly uttered premisses’, then both (38) and (39) should be equally felicitous. By contrast, my account provides a unified account of the distribution of ‘must’ in both deductive and non-deductive contexts. (A referee helpfully suggests that perhaps Giannakidou and Mari and Goodhue have in mind a different characterization of ‘deductive’, on which (38) but not (39) counts as a deductive environment. This may be; but then my point is simply that that characterization needs to be given explicitly. Either way, the point is the same: we need a theory of when ‘must’ is warranted in (what I call) deductive environments.) Goodhue adduces cross-linguistic data which he claims suggest that some languages distinguish a deductive from an epistemic ‘must’, and more work is clearly needed here. It should be noted that distinguishing a deductive from epistemic ‘must’ is of course consistent with my view.

²³As noted above, the ability of our account to deal with these cases very much depends on being able to spell out *Non-Redundancy* in a way that avoids the conclusion that every necessary truth follows in a mutually obvious way from every other one. This conclusion is obviously false, but it is notoriously difficult to give a formal model of this, and I do not attempt to do so here. As far as I can tell, the distribution of ‘must’ in logical and mathematical environments is unexplained by von Fintel and Gillies 2010: since logical and mathematical truths are entailed by anything, our evidence for them will, in von Fintel and Gillies’ terms, always count as “direct”, and thus ‘must’ should never be licensed. As Goodhue (2017) writes, ‘deduction contexts do not provide clear evidence in favor of one account of epistemic modality over another since they are puzzling from the perspectives of both the epistemic account and the indirectness account’. I think this is correct given extant accounts of *Indirectness*; by contrast, the present account of *Indirectness* provides a straightforward and unified explanation of these cases.

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. Consider first 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 *p*. 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 (40):

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

(40) may be misleading, but it strikes me as perfectly felicitous, despite the fact that Mary has direct evidence that there is a party. This is precisely our prediction, since *Strongest Evidence* seems to be suspended here.

It is somewhat harder to find similar cases in which *Non-Redundancy* is suspended, since *Non-Redundancy* already has an element of context-sensitivity built in (since what counts as mutually obvious in the relevant sense is 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. One candidate context is the sort that we find when a speaker wishes to signal that her interlocutor is being dense, and thus spells out reasoning even when it is mutually obvious:

- (41) a. I can't believe it. Are you really firing me?
b. Well, let's see. I'm your boss. Under your contract, I have complete discretion to terminate your employment at any time. I told you in writing that you're fired. So, yep, I'm firing you.

In contexts like this, the boss is spelling out reasoning that is mutually obvious; her assertion seems to be felicitous, if obnoxious, and so *Non-Redundancy* seems to be suspended here. We thus predict that *Indirectness* will be suspended as well—that is, that 'must' is perfectly felicitous here, even though the boss's evidence about whether she is firing her employee is perfectly direct. This seems to be borne out: (41) is also felicitous (if obnoxious) if we put a 'must' over the final sentence. These points suggest that, as I have argued, *Indirectness* is indeed a pragmatic inference, which will therefore be suspended when the underlying pragmatic principles are suspended.

More generally, the present discussion shows that deriving *Indirectness* in the way I have proposed not only gives us an account of where it comes from, but also gives us a *predictive* account, one which allows us to make sense of the subtle pattern of judgments about the felicity

of ‘must’-claims by reducing them to judgments about redundancy.

4 Support

Support plus independently motivated pragmatic principles thus provides a theoretically and empirically satisfying explanation of *Indirectness*. I turn now to the question of how to account for *Support*. I briefly survey and criticize the few extant proposals before sketching my own proposal. (Just as my account of *Indirectness* in the last section was separable from my empirical claims in §2, likewise my account of *Support* in this section is separable from my account of *Indirectness* and empirical claims.)

4.1 Negative remarks

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 p’ means roughly ‘This argument entails p’, 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 needs to be saturated by an argument made salient by context.

A solution along these lines, natural though it is, does not work. The issue is that, as we saw above, ‘might’ and ‘can’, unlike ‘must’, do not require a salient argument. We have already seen a variety of intuitive and experimental evidence for this in §2; to recall the contrast, consider (42):

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

If the conversation ends here, then, as *Support* predicts, the ‘must’ variant in (42-c) is infelicitous as it stands, without an argument. By contrast, the ‘might’ variant in (42-a), like the non-modal variant in (42-b), is acceptable without an argument. In short: ‘might’, unlike ‘must’, does not require an argument for felicity.

If we take Stone’s approach, however, then, assuming 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 stipulate that ‘might’ does not have a lexical argument place for an argument. But going this way would lead to a serious new issue. Assuming we treat ‘cannot’ as ‘not (can)’, and that epistemic ‘can’ is equivalent to epistemic ‘might’, if we go this way, we will predict that ‘cannot’, like ‘can’, does not have an anaphoric requirement for an argument any more than (unembedded) ‘can’ 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’. Thus, for instance, consider (43):

- (43) [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.

Just as for ‘must’, and just as predicted by *Support*, if the exchange ends here, the ‘can’t’ variant without an argument in (43-a) is marked. As for ‘must’, then, ‘can’t’ requires that an argument be made salient. If we treat ‘might’ and ‘can’ as lacking a lexical requirement for an argument, then we will also predict that ‘can’t’ lacks a lexical requirement for an argument, and thus we will fail to make sense of the fact that ‘can’t’ requires an argument.

One might at this point simply claim that ‘can’ and ‘might’, which I have so far assumed are semantically identical, in fact differ in just the respect that ‘can’ has an argument place for an argument, while ‘might’ does not. This claim looks plausible enough in English, since ‘can’ is rarely interpreted epistemically when it is not negated, and ‘might’ cannot be easily negated; we could maintain that these differ precisely in the respect suggested. However, the cross-linguistic picture suggests that this line will be hard to maintain. Many languages use the same epistemic possibility modal in positive and negative contexts, like German (‘können’) and French (‘se pouvoir’). In those languages, informants tell me (based on contrasts like those given above) that we find the same alternation as in English: matrix possibility modals in sentences like (44) and (45) do not require a salient argument, whereas negated possibility modals (hence with a necessity meaning), as in (46) and (47), do require a salient argument (just as predicted by *Support*).

- (44) Der Mörder könnte der Gärtner sein.
The murderer might the gardener be.
‘The murderer might be the gardener’.

- (45) Il se peut que le meurtrier soit le jardinier.
It might be that the murderer is the gardener.

‘The murderer might be the gardener’.

(46) Der Mörder könnte nicht der Gärtner sein.
The murderer might not the gardener be.
‘The murderer can’t be the gardener’.

(47) Il ne se peut pas que le meurtrier soit le jardinier.
It not might be that the murderer is the gardener.
‘The murderer can’t be the gardener’.

To pursue the line under discussion beyond English, we would have to maintain that there are two words ‘können’ in German, one which only appears in positive contexts and has no requirement for a salient argument, and one which appears only in negative contexts and requires a salient argument. Likewise for French, *mutatis mutandis*, and presumably for many other languages. Such an approach is obviously unattractive.²⁴

4.2 *Support as a manner implicature*

We can avoid these problems by deriving *Support* as a manner implicature. The derivation depends on the assumption that the interpretation of unembedded epistemic modals goes by way of the common evidence of the interlocutors. This is not a commitment about the *semantic value* of epistemic modals (and thus about their interpretation when embedded), and so this assumption is compatible with a variety of semantics for epistemic modals, in particular the standard modal semantics of Kripke 1963; Hintikka 1962; Kratzer 1977, 1981.²⁵ On those approaches, an unembedded assertion of ‘Must p ’ is interpreted as meaning ‘ p is entailed by the contextually relevant information’. The present claim can then be interpreted as a precisification of these approaches: that is, as the claim that the contextually relevant information for unembedded modal claims is typically determined by the *common evidence* of the group of interlocutors. As Stalnaker 2014; Mandelkern 2017a, 2018 argue, this kind of approach provides a way to make sense of the basic update effects of modal claims, in particular the way they are used to coordinate on the structural properties of the common ground. In particular, this approach predicts that an assertion of ‘Must p ’ is felt to be a proposal to make p entailed by the common ground, since, if it is common ground that the common evidence entails p , then, given plausible assumptions about the relation between evidence and belief, p itself will be common ground. This approach thus validates *Pragmatic Strength*—an assertion of ‘Must p ’,

²⁴Similar criticisms extend to any kind of presuppositional account of *Support*, since presuppositions project through negation; as well as 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 (see Moss 2015 for a similar proposal). As far as I can tell, that approach faces the same dilemma.

²⁵Setting aside the role of a normality ordering source in Kratzer’s semantics, the source of its characteristic weakness.

if accepted, has the effect of adding p to the common ground.²⁶ Likewise, given mild assumptions about the connection between the common ground and common evidence, it will follow that if a group accepts p (i.e. if the common ground entails p), then they also accept that the common evidence entails p , and thus they also accept $\ulcorner \text{Must } p \urcorner$.

The key feature of this approach for present purposes is that an assertion of $\ulcorner \text{Must } p \urcorner$ ends up having the same update effect as an assertion of p alone: if a group accepts p , they accept $\ulcorner \text{Must } p \urcorner$; and if they accept $\ulcorner \text{Must } p \urcorner$, they accept p . So $\ulcorner \text{Must } p \urcorner$ and p , on this account, are informationally equivalent in the sense that any context which entails one entails the other. Note, moreover, that p is structurally simpler than $\ulcorner \text{Must } p \urcorner$.²⁷ Since they have the same basic update effect, they will therefore be in competition: if a speaker chooses a more complex way to achieve the same basic effect, this choice will require some explanation. What could explain this choice? In the present framework, although $\ulcorner \text{Must } p \urcorner$ and p are informationally equivalent, they have *different subject matters*. p on its own is simply about whether p is true. $\ulcorner \text{Must } p \urcorner$, by contrast, is about whether p follows from the interlocutors' common evidence.²⁸ An assertion of $\ulcorner \text{Must } p \urcorner$ is thus like an assertion of p in that it is a proposal to commonly accept p ; but unlike an assertion of p in that it is a proposal to accept p *which goes by way of calling attention to the interlocutors' common evidence for p* . Given that the speaker will always have a choice between a more complex alternative which calls attention to common evidence in this way, and a less complex alternative which does not, an assertion of $\ulcorner \text{Must } p \urcorner$ will be felt to be, not just a proposal to accept p , but a proposal to accept p *with special attention to the common evidence for p* .

What kind of evidence would be worth drawing attention to in this way? One possibility is that the evidence is just that the speaker herself endorses p . But there is nothing noteworthy about this kind of evidence: this kind of endorsement is in play *whenever* the speaker proposes an update with p —including when she does so by asserting p alone—and thus, if this were the only relevant common evidence, it would not explain why the speaker did not just assert p , the structurally simpler alternative. Calling attention to the speakers' evidence for p will thus make sense only if the speaker wants her interlocutors to accept p based on *substantive evidence* for p —evidence that goes beyond the speaker's endorsement of p . And if the speaker wants her interlocutors to accept p based on evidence like this, then she will, of course, need to ensure that such evidence is available to them—either by providing it explicitly, or by otherwise ensuring

²⁶Though it does not commit us to the stronger, and more controversial, hypotheses mentioned above (that $\ulcorner \text{Must } p \urcorner$ entails p or entails knowledge or certainty of p). The former of these would follow, though, if we treat evidence as a factive notion.

²⁷There is controversy about how to spell out the notion of (relevant) relative structural complexity, but this claim seems uncontroversial to me; it will follow e.g. on the account given in Katzir 2007, according to which alternatives are calculated by the deletion or replacement of nodes at LF. Likewise $\ulcorner \text{Not } p \urcorner$ will be a relevant alternative to $\ulcorner \text{Can't } p \urcorner$.

²⁸There are a variety of ways of spelling out this intuition; a simple approach just says that the subject matter of an assertion of p —the question raised by such an assertion—is simply the question $\{p, \bar{p}\}$. See Lewis 1988; van Kuppevelt 1995; Ginzburg 1995a,b; Roberts 2012; Yablo 2014; Bledin and Rawlins 2016 and citations therein.

that her interlocutors can recover it from the common ground.

Let me pause to note the importance of this last point. It may look like, to explain *Support*, all we have to do is endorse something like the above hypothesis about ‘must’, on which ‘Must p’ says that the common evidence entails *p*. But this leaves unexplained why the evidence in question must be *substantive* in the sense just sketched, since *whenever* a speaker proposes to update the common ground with *p*—either by asserting ‘Must p’ or by asserting *p*—she is creating common evidence that supports *p* by publicly endorsing *p*. ‘Must p’ requires something more than this: again, it requires that the speaker ensure that a *substantive* argument for *p* has been made available. The conclusion that such an argument is required is only obtained when we compare ‘Must p’ with *p*. Once we regard these as being in competition, we can reason that, by choosing the more complex variant, the speaker is drawing attention to substantive common evidence for *p*.

Summarizing this reasoning:

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

Importantly, this approach to *Support* avoids the objection to existing views given above. In particular, this approach predicts that all and only strong epistemic necessity modals will require support. The reason for this is that a key step in the reasoning above was the assumption that ‘Must p’ and *p* alone are in competition. That, in turn, follows because of the assumption

that ‘Must p’ and p are informationally equivalent. This assumption will be merited for all strong epistemic necessity modals (in the case of ‘Can’t p’, the competition is of course with ‘Not p’), but *only* for strong epistemic necessity modals: weaker epistemic modal claims uncontroversially have a weaker update effect than p alone. Thus, unlike existing approaches, we rightly predict that strong epistemic necessity claims (‘Must p’, ‘Can’t p’, and so on) require an argument, while weaker epistemic modal claims (‘Might p’, etc.), like non-modal claims, do not.

Is *Support* ever cancelled in the way that, as we saw above, *Indirectness* can be? It is not clear to me that it is. The pragmatic principle underlying the derivation is something like ‘If p and q are alternatives with the same basic update effect, use the structurally simpler one unless you have a good reason not to’. It is not clear that this principle is ever suspended.²⁹ On the other hand, a derivation of *Support* as a manner implicature predicts that a violation of *Support*—failing to provide an argument with a ‘must’-claim—will be relatively innocuous: it will lead to puzzlement about the speaker’s choice of ‘Must p’ rather than p, but not to total communicative breakdown. This prediction seems to be in line with intuition and the experimental results reported above.

Let me close this section by highlighting a prediction of the present proposal: any construction which has the features which played an essential role in the derivation is predicted to give rise to a *Support*-like constraint. That is, any construction which has p as a structurally simpler relevant alternative; which has the same basic update effect as an assertion of p; and which highlights the speakers’ collective doxastic relation to p, is, *ceteris paribus*, predicted to be degraded without a salient argument for its complement (and thus also to give rise to *Indirectness*, via the pragmatic reasoning discussed in the last section). Further research should explore whether this prediction is borne out.^{30,31}

5 Weakness

This concludes the core of my response to Karttunen’s Problem. Before concluding, I take up two residual questions. First, in this section, I address how my account explains the felt weakness of some ‘must’-claims, and some related further contrasts between ‘must’-claims and non-modal claims.

²⁹See Lauer 2013, Ch. 9 for a defense of the view that some implicatures are neither optional nor cancelable (what he calls ‘Need a Reason’ implicatures).

³⁰Superficial exploration suggests it is; thus for instance ‘Our evidence entails that p’, ‘Let’s agree that p’, and ‘It is clear that p’ all seem to have these properties, and do indeed appear to be governed by corollaries of *Support* and *Indirectness*. See Barker 2009 for related discussion; thanks to Kai von Stechow for discussion.

³¹See Degen et al. 2015 for a different manner-implicature based approach to explaining the behavior of ‘must’ which, like the present approach, relies on the assumption that ‘Must p’ is a costly alternative to p whose use must be somehow explained, but which focuses on *Indirectness*.

5.1 Uncertain ‘must’

There is some sense in which ‘must’ can be felt to weaken a claim. For instance, as Lassiter (2016) observes, agents are sometimes happy to assert sentences along the lines of ‘Must p , but I don’t know p for sure’, as in Lassiter’s naturalistic examples (48) and (49):

- (48) This is not non-stick. There must be others but I don’t know for sure.
- (49) This is a very early, very correct Mustang that has been in a private collection for a long time. . . The speedo[meter] shows 38,000 miles and it must be 138,000, but I don’t know for sure.

Sentences like (48) and (49) sound somewhat marked, but they do seem preferable to the corresponding non-modal variants:

- (50) ??This is not non-stick. There are others but I don’t know [that] for sure.

A related observation is due to Degen et al. (2016), who use a series of experiments to show that participants tend to assert ‘Must p ’ more when the speaker has weaker evidence for p , and to assert p more when she has stronger evidence for p . Likewise, listeners tend to infer that speakers of ‘Must p ’ have weaker evidence for p than speakers of p alone.

How can we make sense of this sense that ‘must’ is weak, particularly in light of the data in §3.1 which suggest that ‘must’ is pragmatically strong? One option, of course, would be to take these data to show that, after all, ‘must’ is not pragmatically strong. But this would, of course, leave unexplained the data that motivated *Pragmatic Strength* above. And it is not clear that the data here are inconsistent with *Pragmatic Strength*. They would be if, for instance, the second conjunct of (48) had to be interpreted as a proposal to leave it open that there are not other non-stick pans. But (48) could instead be a proposal to update with the proposition that there are other non-stick pans, *and* to record that the speaker’s evidence for this proposition is less certain than it could be. Likewise for (49). Indeed, as Paul Portner (p.c.) notes, examples like (48) and (49) have the sense of cases in which the speaker’s is setting aside her mental state as being not all that relevant: it is natural to interpret them as claims in which the speaker proposes some update (that there are other non-stick pans, that the Mustang has 138,000 miles on it), while also acknowledging that the her evidence for the claim is less certain than it could be. So, in light of these considerations and the data reviewed in §3.1, it seems more promising to me to try to make sense of the intuition of the weakness of ‘must’ while holding onto *Pragmatic Strength*.

And indeed, it seems to me that there are at least two ways in which we can do so in the context of the account I have given. The first is based on the simple observation that, on my account—as on others—‘must’ is appropriate only when the speaker’s evidence is indirect. Indirect evidence is not always weaker than direct evidence (in logical contexts, it is not); but

it often is, and this may go much of the way towards accounting for the intuition that ‘must’ is often associated with weaker evidence.

The second is based on the observation that, on my account, in uttering ‘Must p ’, a speaker proposes that her interlocutors accept p , but she proposes that they do so on the basis of an argument that is publicly available, *not* on the basis of trust that she has private evidence for p . And asking one’s interlocutors to accept p on the basis of a shared argument, rather than on the basis of one’s own epistemic authority, will plausibly be permissible when one’s own epistemic standing towards p is something short of what would be required for an outright assertion of p alone. This is because, when the speaker proposes accepting p on the basis of a shared argument rather than her own epistemic authority, the whole group is provided the evidence in question, and tasked with determining whether or not the evidence warrants the conclusion of p —rather than simply being asked to accept the speaker’s authority. Taking this tack may well be appropriate when the speaker thinks that her evidence warrants accepting p , but is not entirely sure, and so wants the whole group to decide whether or not p is warranted on the basis of that evidence. By sharing her evidence for p with the group, she allows the group to decide whether *that* evidence warrants accepting p —bringing closer scrutiny, from the whole group of deliberators, to the decision. Such a move will generally be warranted only when she thinks the decision merits such scrutiny; in some cases (though, of course, not all), this will be because the evidence in question provides her with something short of certain knowledge of p . In those cases, sharing your evidence for p allows the group to keep track of the inferential path that led them to accept p . This is especially useful to do when you think there is a chance you will later discover that p is false, since belief revision is most efficient when you have kept track of the inferential relations between your beliefs.³² In short, then, because on my account ‘Must p ’ amounts to a proposal to accept p based on a shared argument, asserting ‘Must p ’ thus will be compatible with having weaker evidence for p than is required for asserting p alone. This helps explain Lassiter’s observation that ‘Must p and I don’t know p for sure’ is improved as compared with the non-modal variant.

Let me note that this reasoning assumes that there is some difference between *knowledge* and *certain knowledge*. In order to account for the data I used to motivate *Pragmatic Strength* in §3.1, we must maintain that you need to know p in order to assert ‘Must p ’. And, as Lassiter observes, there is a contrast between ‘Must p , but I don’t know p [for sure/with certainty]’, which is perhaps marginally acceptable, versus ‘Must p , but I don’t know p ’, which, as we have seen, is totally unacceptable. This contrast suggests that there is indeed a difference between certain knowledge and knowledge (though I won’t try to give an account of the difference here; see Moss 2018, §5.9 for some discussion). Further evidence for this comes from the felicity of exchanges like (51):

(51) a. [A:] I know whether John is coming to the party.

³²In other words, when you have a sophisticated system of epistemic entrenchment; see Hansson 2016 for an overview of relevant literature.

- b. [B:] Ok, but do you know for sure?
- c. [B:] Ok, but are you certain?

5.2 Obligatory ‘must’

This helps explain why ‘Must p ’ is sometimes associated with relatively weak evidence for p : ‘Must p ’ requires sharing that evidence and requires that evidence to be indirect, both of which moves, for different reasons, are often associated with weaker evidence for p .

But what we have said so far leaves unexplained a related issue: why is an assertion of p alone sometimes unacceptable in a parallel context, *even when the speaker’s evidence for p is shared*? The point is brought out clearly by the following case, adapted slightly from Ninan 2014:

- (52) [A and B are friends with a couple, Carl and Diane. It is common knowledge for A and B that A and B have been dating for a long time and are likely to get married at some point in the future. Suppose that, prior to the following dialogue, B has not heard any recent news concerning Carl and Diane’s relationship. Now consider:]
- a. [A:] Carl proposed to Diane yesterday!
 - b. [B:] At last! She must have said ‘yes’.
 - c. [B:] ?At last! She said ‘yes’.

(52-c) is somewhat marked in comparison with (52-b): it looks like we have a preference for using the ‘must’ variant here. Intuitively, this is because ‘She must have said ‘yes’’ makes clear that B’s conclusion is drawn from the evidence in common, whereas ‘She said ‘yes’’ makes it unclear on what basis B drew her conclusion.

We have already seen that ‘must’-claims draw attention to commonly available evidence for the prejacent in question, and so we have a good story about why the ‘must’ variant in (52-b) sounds acceptable in this case. On our account, using ‘must’ makes B’s inferential path explicit: by using the ‘must’ variant, B is proposing to accept that Diane said ‘yes’ *based on the evidence the speakers have in common about the strength of Carl and Diane’s relationship*. And B’s evidence in this case is weak enough that she has every reason to think that, although she and her interlocutors now have reason to accept that Diane said ‘yes’, they may later have reason to revise this commitment, and so it is crucial that they keep track of the inferential path that led to this acceptance.

But why does the non-modal variant sound marked here? The reason is that using the non-modal variant leaves it open that the speaker has evidence that Diane said ‘yes’ *which goes beyond what is shared in the context*, and that it is based on this evidence, rather than their shared evidence, that B is proposing to update with the claim that Diane said ‘yes’. In other words, even though A and B do have shared evidence here in support of this claim, only ‘must’ makes it clear that the update is *based on that shared evidence*.

To illustrate the point, suppose that A and B discover that Carl and Diane have recently been fighting a lot. Should A revise her belief that Diane said ‘yes’? Intuitively, she should. And if B uses the ‘must’ variant, then A will know that she should, since she will know the basis of B’s conclusion that Diane said ‘yes’—she will know that B was just making an inferential leap from their shared evidence. But not so if B uses the non-modal variant without further elaboration, in which case A will not know on what basis B concluded that Diane said ‘yes’, and thus will not know whether revision is merited. (Things will be otherwise, of course, if B makes her reasoning explicit, even with the non-modal variant; and, as predicted, in that case the non-modal claim is substantially improved, as in: ‘At last! She said ‘yes’, I can’t imagine her turning him down.’)

In short: in this case, given our account of ‘must’, (52-b) amounts to a proposal to accept that Diane said ‘yes’ on the basis of the evidence which is already shared; whereas (52-c) amounts to a proposal to accept that Diane said ‘yes’ on the basis of the speaker’s epistemic authority alone. In this case, taking the latter option obscures the inferential path which lead to B’s conclusion—an inferential path which may well be crucial for A to have access to, given the defeasibility of B’s inference that Diane said ‘yes’.³³

5.3 Violations of *Non-Redundancy*

A referee for this journal points out that there are cases that are something like the converse of Ninan’s case: cases in which a non-modal claim seems to be felicitous, but the corresponding ‘must’-claim is, at least *prima facie*, infelicitous. To see the point, suppose first that A and B are looking out at the rain, and it is common ground that both can see that it is raining. Then it may be acceptable for A to say:

(53) It’s raining.

This on its own is *prima facie* surprising for our account. If *Common Ground Settlement* is right, then the common ground already entails that it is raining, and so A should not be able to say this. But it was of course already recognized by Stalnaker 1974 that constraints on redundant assertions are defeasible: they may be suspended when the purpose of an assertion is not to communicate an increment of information, but rather to make idle conversation, set the topic for a conversation, raise a certain matter to salience, and so on. Thus we could imagine A saying (53) as a way of setting up the question of whether the picnic will be cancelled:

(54) It’s raining. Do you think the picnic will be canceled?

³³This explanation may extend to another class of cases, brought to my attention by Agnes Callard, involving moral judgments. ‘Lee was late, and therefore disrespectful’ is felicitous, but ‘Lee was late, and therefore must be disrespectful’ is strange. The explanation, again, may be that we expect speakers to put the full weight of their epistemic authority behind assertions like this.

There is a contrast here, however, with the corresponding ‘must’-variant, which is quite odd:

(55) It must be raining. Do you think the picnic will be canceled?

The considerations adduced above to account for Ninan’s case can also help account for this contrast. On our account, again, ‘Must p ’ draws attention to the interlocutors’ shared evidence for p . If A is using p just as a conversation starter—to set the topic for the conversation, raise p to salience, etc.—there will be no reason to use a variant which focuses particular attention on the interlocutors’ shared evidence for p , and so using ‘Must p ’ over p will strike us as strange: the use of a more complex variant to communicate the same basic information calls for an explanation, and, when we do not have any reason to draw attention to our evidence in support of p , such an explanation will be lacking. In other words, (55) is strange for the same reason that (56) is in this context:

(56) Based on our visual perception, it’s raining. Do you think the picnic will be canceled?

This account yields a striking prediction: in a variant of the case just given in which the interlocutors *do* wish to call attention to their evidence, the ‘must’-variant should be acceptable, and indeed preferred. This prediction seems to be borne out. Consider the following variant on von Fintel and Gillies (2010)’s epistemologist case (34). Suppose that A and B are epistemologists. Both accept that visual evidence can directly supply knowledge. So both accept that, from the visual perception that it is raining, they can immediately conclude that it is raining: this conclusion is mutually obvious to them. Nonetheless, A wants to start a conversation about the precise nature of this inference. Suppose further that A and B are standing at the window, looking out at the rain. A can say:

(57) It must be raining. In other words, our visual perception shows us that it is raining, and that suffices for us to know that it’s raining. But how exactly does our visual perception license this inference?

By contrast, the non-modal variant here is at least somewhat dispreferred. Suppose A said instead:

(58) ?It’s raining. In other words, our visual perception shows us that it is raining, and that suffices for us to know that it’s raining. But how exactly does our visual perception license this inference?

In this non-modal variant, ‘in other words’ seems like a non-sequitur. (This judgment is a bit subtle—after all, ‘it’s raining’ could also be used to set the topic of rain-evidence. But the ‘must’ certainly seems preferred here.)

Note that this case is subtly different from (34). There, the inference from visual perception to knowledge is being called into question. In the present case, this inference is not in doubt;

it is simply the subject of discussion. So this is not a case in which this inference ceases to be mutually obvious, as in the case of (34); instead, this is a case in which this inference remains mutually obvious, but is the subject of discussion, and so a ‘must’ is merited.

In sum, then, *both* ‘must’-claims and non-modal claims can be used when there is a shared argument from which p follows in a mutually obvious way. In such cases, *Non-Redundancy/ Common Ground Settlement* can be suspended if there is a clear purposes served in doing so. If the assertion is being used to set p as the topic of conversation, draw attention to p , etc., then the ‘must’-variant is clearly dispreferred; by contrast, if the assertion is being used (at least in part) to call attention to the interlocutors’ evidence for p , then the ‘must’-variant is preferred. This contrast is predicted by our account, on which $\lceil \text{Must } p \rceil$ and p are in general in competition with each other, and speakers use the former only if they wish to update with p while drawing attention to their evidence for p .

6 Embeddings

The final question I address concerns how to extend the present solution to Karttunen’s Problem to embeddings of ‘must’. First, note that *Indirectness* persists when ‘must’-claims are embedded under connectives as in (59-a) and (59-b), which communicate that the speaker’s evidence that Mark came to the party is suitably indirect:

- (59) [Who came to the party?]
- a. John did, and Mark must have as well; they’re inseparable.
 - b. John did, and Mark must have as well.
 - c. John did and Mark did.

Support also seems to be in play in cases like this. It is satisfied in (59-a), which is felicitous as it stands. (59-b) is felicitous only if we can accommodate an argument that Mark came (for instance, that John and Mark always go out together). By contrast, (59-c) requires no such inference: there need be no salient argument that Mark came to the party.

A natural first response to cases like this is that, since these signals persist through embeddings, they must be semantically encoded. But this would be too quick. On the one hand, it is not at all clear how to capture these data semantically. As Ippolito (2017) shows, encoding *Indirectness* as a semantic presupposition runs into serious difficulty; as I discussed in §4.1, the same goes for *Support*.³⁴ On the other hand, I am hopeful that we can explain these cases pragmatically by generalizing the strategy from the unembedded cases. The details of this go

³⁴Encoding either of them as conventional implicatures would not do any better, since it would predict unconditional projection, and thus face the same counterexamples that make trouble for the presuppositional approaches discussed above, at least on the approach due to Potts 2005, though others have resisted the claim that conventional implicatures project unconditionally (e.g. Martin 2016). And it is clear that they are not part of the main content of ‘must’, since e.g. $\lceil \text{Not must } p \rceil$ clearly does not mean the same thing as $\lceil \text{My evidence for } p \text{ is direct, or fails to entail } p \rceil$.

beyond the scope of this paper, but the basic idea is that the principles that played a crucial role in our derivation, in particular *Non-Redundancy*, hold not only at the level of sentences, but also at a local, sub-sentential level.³⁵ We can then reason roughly as above to derive both *Support* and *Indirectness*; likewise for other connectives.³⁶

‘Must’ under attitudes, as in (60), has a slightly different empirical and theoretical profile:

(60) James thinks Bob must be in his office.

Here, again, the ‘must’ seems to suggest that James’ evidence is indirect, but no correlate of *Support* seems to apply here. How to predict this will depend on the interaction between attitude verbs and ‘must’. The most natural way to treat epistemic modals under attitudes (as in Stephenson 2007b,a) is to take the agent of the attitude predicate to “anchor” the modal base, so that ‘must’ is interpreted as quantifying over the set of worlds compatible with a relevant body of the agent’s evidence. Thus an attitude ascription like (60) says something like: James believes his evidence entails that Bob is in his office.³⁷ It seems to me that there is a natural step from these truth conditions to the conclusion that James’s evidence is indirect—it is only in that case that he would bother to reflect on it at all—though I will not try to spell this out in detail here.

Much more needs to be said on the present topic, then; but this discussion should suffice to give a sense of how the present account can be extended to make sense of the distribution, not only of unembedded ‘must’-claims, but also of embedded ‘must’.

7 Conclusion

The argument of this paper came in three stages. I began by arguing that, in addition to *Indirectness*, we need *Support* to fully characterize the differences in felicity conditions between an assertion of ‘Must p’ and an assertion of p alone (as well as the difference between ‘Must p’ and epistemic modal claims of other strengths). Next, I argued that we can account for *Indirectness* by deriving it pragmatically from *Support*. Finally, I argued that we should account for *Support* as a manner implicature.

The three main parts of this argument are, to a degree, independent: the second and third part depend on the first, but not *vice versa*; and the second and third part are independent of each other. 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

³⁵Such a theory is needed to explain the infelicity of sentences like ‘If John is in Paris, then John is in Paris and he’s having fun’. See e.g. Schlenker 2009, Katzir and Singh 2013, Mayr and Romoli 2016.

³⁶Though importantly this approach will *not* predict that ‘must’ under negation requires an argument, since ‘Not must p’ and ‘Not p’ do not have the same basic update effect, which was an important assumption in our derivation

³⁷Things are a bit more complicated than this, for reasons Yalcin 2007 points out, but need not concern us at present. See discussion in Mandelkern 2019.

conditions between an assertion of ‘Must p’ and an assertion of p alone. In short: because of its more complex form and its difference in subject matter from p alone, an assertion of ‘Must p’, unlike an assertion of p alone, requires that an argument be given for p. From this requirement, in turn, we can conclude that the speaker’s evidence for p is indirect (in a relevant sense), on the basis of general considerations about the flow of information in conversation. I have argued that this approach provides an explanatorily and empirically satisfying account of *Support*; of the subtle pattern of judgments underlying *Indirectness*; and, finally, of the sense in which ‘must’ is felt to be weak, in spite of its pragmatic strength.

I close with an abstract point about the architecture of semantic and pragmatic theories. My proposal rests on the assumption that an assertion of ‘Must p’ and an assertion of p have the same basic update effect, but different semantic values. Indeed, on the semantics I have sketched, ‘Must p’ and p informationally entail one another—in the sense that a context which entails 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 capturing the subtle divergence in felicity conditions between ‘Must p’ and p alone; and so, if we are to adopt an approach like this, then we must distinguish semantic content from pragmatic update effect in theorizing about natural language.

References

- Aikhenvald, A. (2004). *Evidentiality*. Oxford University Press, Oxford.
- Barker, C. (2009). Clarity and the grammar of skepticism. *Mind and Language*, 24(3):253–73.
- Barker, C. and Taranto, G. (2003). The paradox of asserting clarity. In Koskinen, P., editor, *Western Conference on Linguistics (WECOL)*, pages 10–21.
- Bledin, J. and Rawlins, K. (2016). Epistemic resistance moves. In Moroney, M., Little, C.-R., Collard, J., and Burgdorf, D., editors, *Semantics and Linguistic Theory (SALT)*, volume 26, pages 620–640.
- Bronnikov, G. (2008). The paradox of clarity: Defending the missing inference theory. In Friedman, T. and Ito, S., editors, *Semantics and Linguistic Theory (SALT)*, volume 18, pages 144–157, Ithaca, NY. Cornell University.
- Crone, P. (2015). Asserting clarity & managing awareness. In Bade, N., Berezovskaya, P., and Schöller, A., editors, *Sinn und Bedeutung*, volume 20, pages 162–179.
- Degen, J., Kao, J. T., Scontras, G., and Goodman, N. D. (2015). A cost- and information-based account of epistemic *must*. Poster at 28th Annual CUNY Conference on Human Sentence Processing.
- Degen, J., Scontras, G., Trotzke, A., and Wittenberg, E. (2016). Definitely, maybe: Approaching speaker commitment experimentally. Manuscript, Stanford University.
- Faller, M. (2001). Remarks on evidential hierarchies. In Beaver, D., Kaufmann, S., Clark, B., and Casillas, L., editors, *Proceedings of the ‘Semfest’*. CSLI Publications: Stanford, CA.
- Faller, M. (2002). *Semantics and Pragmatics of Evidentials in Cuzco Quechua*. PhD thesis, Stanford University, Stanford, California.
- Faller, M. (2012). Evidential scalar implicatures. *Linguistics and Philosophy*, 35(4):285–312.
- von Stechow, K. and Gillies, A. (2010). Must...stay...strong! *Natural Language Semantics*, 18(4):351–383.
- von Stechow, K. and Gillies, A. (2018). Still going strong. <https://semanticsarchive.net/Archive/GMzZTM0Y/>.
- Franke, M. and de Jager, T. (2007). The relevance of awareness. In Aloni, M., Dekker, P., and Roelofsen, F., editors, *The Sixteenth Amsterdam Colloquium*, pages 91–96.

- Giannakidou, A. and Mari, A. (2016). Epistemic future and epistemic MUST: nonveridicality, evidence, and partial knowledge. In Blaszack, J., Giannakidou, A., Klimek-Jankowska, D., and Mygdalski, K., editors, *Mood, Aspect and Modality: What is a Linguistic Category?* University of Chicago Press.
- Giannakidou, A. and Mari, A. (2017). A unified analysis of the future as epistemic modality: The view from Greek and Italian. *Natural Language and Linguistic Theory*, pages 1–45.
- Gillies, A. (2001). A new solution to Moore’s paradox. *Philosophical Studies*, 105(3):237–250.
- Ginzburg, J. (1995a). Resolving questions. i. *Linguistics and Philosophy*, 18(5):459–527.
- Ginzburg, J. (1995b). Resolving questions. ii. *Linguistics and Philosophy*, 18(6):567–609.
- Goodhue, D. (2017). *Must* φ is felicitous only if φ is not known. *Semantics & Pragmatics*, 10(14).
- Grice, P. (1989). *Studies in the Way of Words*. Harvard.
- Groenendijk, J., Stokhof, M., and Veltman, F. (1996). Coreference and modality. In *Handbook of Contemporary Semantic Theory*, pages 179–216. Oxford: Blackwell.
- Hansson, S. O. (2016). Logic of belief revision. In Zalta, E. Z., editor, *The Stanford Encyclopedia of Philosophy*, Winter 2016 Edition. <https://plato.stanford.edu/archives/win2016/entries/logic-belief-revision/>.
- Hintikka, J. (1962). *Knowledge and Belief: An Introduction to the Logic of Two Notions*. Cornell University Press, Ithaca, NY.
- Ippolito, M. (2017). Constraints on the embeddability of epistemic modals. To appear in *Sinn und Bedeutung 21*.
- Karttunen, L. (1972). Possible and must. In Kimball, J., editor, *Syntax and Semantics*, volume 1, pages 1–20. Academic Press, New York.
- Katzir, R. (2007). Structurally-defined alternatives. *Linguistics and Philosophy*, 30(6):669–690.
- Katzir, R. and Singh, R. (2013). Hurford disjunctions: embedded exhaustification and structural economy. In Etxeberria, U., Fălăuş, A., Irurtzun, A., and Leferman, B., editors, *Sinn und Bedeutung*, volume 18.
- Kratzer, A. (1977). What ‘must’ and ‘can’ must and can mean. *Linguistics and Philosophy*, 1(3):337–355.
- Kratzer, A. (1981). The notional category of modality. In Eikmeyer, H. and Rieser, H., editors, *Words, Worlds, and Contexts: New Approaches in Word Semantics*, pages 38–74. de Gruyter.
- Kratzer, A. (2012a). *Modals and Conditionals*. Oxford University Press, Oxford.
- Kratzer, A. (2012b). What ‘must’ and ‘can’ must and can mean. In *Modals and Conditionals*, pages 1–20. Oxford University Press.
- Kripke, S. (1963). Semantical considerations on modal logic. *Acta Philosophica Fennica*, 16:83–94.
- Kripke, S. (2009). Presupposition and anaphora: Remarks on the formulation of the projection problem. *Linguistic Inquiry*, 40(3):367–386.
- van Kuppevelt, J. (1995). Discourse structure, topicality and questioning. *Journal of Linguistics*, 31(1):109–147.
- Lassiter, D. (2016). *Must*, knowledge, and (in)directness. *Natural Language Semantics*, 24(2):117–163.
- Lauer, S. (2013). *Towards a Dynamic Pragmatics*. PhD thesis, Stanford University, Stanford, California.
- 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.
- MacFarlane, J. (2004). In what sense (if any) is logic normative for thought? Presented at the Central Division APA.
- Mandelkern, M. (2017a). *Coordination in Conversation*. PhD thesis, Massachusetts Institute of Technology.
- Mandelkern, M. (2017b). A solution to Karttunen’s problem. In Truswell, R., Cummins, C., Heycock, C., Rabern, B., and Rohde, H., editors, *Sinn und Bedeutung 21*, pages 827–844.
- Mandelkern, M. (2018). How to do things with modals. *Manuscript*.
- Mandelkern, M. (2019). Bounded modality. *The Philosophical Review*, 181(1).
- Martin, S. (2016). Supplemental update. *Semantics & Pragmatics*, 9(5):1–61.

- Matthewson, L. (2015). Evidential restrictions on epistemic modals. In Alonso-Ovalle, L. and Menendez-Benito, P., editors, *Epistemic Indefinites*. Oxford University Press, New York.
- Mayr, C. and Romoli, J. (2016). A puzzle for theories of redundancy: Exhaustification, incrementality, and the notion of local context. *Semantics and Pragmatics*, 9(7):1–48.
- Montague, R. (1970). Universal grammar. *Theoria*, 36(3):373–398.
- Moore, G. (1942). A reply to my critics. In Schilpp, P., editor, *The Philosophy of G.E. Moore*. Northwestern University, Evanston, IL.
- Moss, S. (2015). On the semantics and pragmatics of epistemic vocabulary. *Semantics and Pragmatics*, 8(5):1–81.
- Moss, S. (2018). *Probabilistic Knowledge*. Oxford University Press.
- Murray, S. (2014). Varieties of update. *Semantics and Pragmatics*, 7(2):1–53.
- Ninan, D. (2014). Taste predicates and the acquaintance inference. In Snider, T., D’Antonio, S., and Weigand, M., editors, *Semantics and Linguistic Theory (SALT)*, volume 24, pages 290–309.
- Ozturk, O. and Papafragou, A. (2015). The acquisition of epistemic modality: From semantic meaning to pragmatic interpretation. *Language Learning and Development*, 11(3):191–214.
- Potts, C. (2005). *The Logic of Conventional Implicatures*. Oxford University Press, Oxford.
- Roberts, C. (1998/2012). Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics*, 5(6):1–69.
- Schlenker, P. (2008). Be articulate: a pragmatic theory of presupposition projection. *Theoretical Linguistics*, 34(3):157–212.
- Schlenker, P. (2009). Local contexts. *Semantics and Pragmatics*, 2(3):1–78.
- Scott, D. (1970). Advice on modal logic. In Lambert, K., editor, *Philosophical Problems in Logic*, pages 143–173. Dordrecht Reidel.
- Sherman, B. (2018). Open questions and epistemic necessity. *The Philosophical Quarterly*.
- Silk, A. (2016). *Discourse contextualism: A framework for contextualist semantics and pragmatics*. Oxford University Press.
- Stalnaker, R. (1970). Pragmatics. *Synthese*, 22(1-2):272–289.
- Stalnaker, R. (1973). Presuppositions. *Journal of Philosophical Logic*, 2(4):447–457.
- Stalnaker, R. (1974). Pragmatic presuppositions. In Munitz, M. K. and Unger, P., editors, *Semantics and Philosophy*, pages 197–213. New York University Press, New York.
- Stalnaker, R. (1978). Assertion. In Cole, P., editor, *Syntax and semantics*, volume 9, pages 315–322. New York: Academic Press.
- Stalnaker, R. (2002). Common ground. *Linguistics and Philosophy*, 25(5-6):701–721.
- Stalnaker, R. (2014). *Context*. Oxford University Press.
- Stephenson, T. (2007a). Judge dependence, epistemic modals, and predicates of personal taste. *Linguistics and Philosophy*, 30(4):487–525.
- Stephenson, T. (2007b). A parallel account of epistemic modals and predicates of personal taste. In Puig-Waldmueller, editor, *Sinn und Bedeutung*, volume 11, pages 583–97.
- Stone, M. (1994). The reference argument of epistemic *must*. In *Proceedings of IWCS 1*, pages 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*. PhD thesis, University of Amsterdam.
- Willett, T. (1988). A cross-linguistic survey of the grammaticalization of evidentiality. *Studies in Language*, 12(1):51–97.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press.
- Wittgenstein, L. (2001/1953). *Philosophical Investigations*. Wiley-Blackwell, 3rd edition.

- Wolf, L. and Cohen, A. (2011). Clarity as objectivized belief. In Égré, P. and Klinedinst, N., editors, *Vagueness and Language Use*, pages 165–190. Palgrave Macmillan UK, London.
- Yablo, S. (2014). *Aboutness*. Princeton University Press.
- Yalcin, S. (2007). Epistemic modals. *Mind*, 116(464):983–1026.