

Must ϕ is felicitous only if ϕ is not known*

Daniel Goodhue
McGill University

Submitted 2016-06-05 / First decision 2016-09-13 / Revision received 2017-04-07 /
Accepted 2017-05-03 / Final version received 2017-05-27 / Published 2017-12-20 /
Final typesetting 2025-01-21

Abstract In recent work (von Fintel & Gillies 2010, Matthewson 2015, Lassiter 2016, Mandelkern 2016), epistemic modals have been claimed to have felicity conditions that require the evidence for the prejacent to be indirect. In contrast, I argue that epistemic modals have felicity conditions that require that the prejacent is not known as claimed in Giannakidou & Mari 2016. New linguistic data is produced in support of this position. The proposed account is argued to explain the new evidence better than accounts that rely on indirectness. The evidence in favor of this account also militates in favor of a weak semantics for *must* ϕ . In light of these findings, future prospects are explored. In particular, I suggest that this proposal paves the way for the felicity conditions of epistemic *must* to be derived as a conversational implicature. Furthermore, I demonstrate that a purported counterexample to the proposal, *must* ϕ statements in the conclusions of deductions, is a problem for indirectness accounts as well, and I suggest a way forward.

Keywords: epistemic modality, knowledge, indirect, evidential, inference, deduction, implicature

-
- * This paper supersedes Goodhue 2016. Thank you to the editor, Michael Franke, and the anonymous reviewers at *S&P*. Thanks to Alan Bale, Brian Buccola, Kai von Fintel, Thony Gillies, Valentine Hacquard, Peter Klecha, Dan Lassiter, Salvador Mascarenhas, Lisa Matthewson, Kyle Rawlins, Dee Ann Reisinger, Raj Singh, Michael Wagner, Matthijs Westera, the audiences at NELS 46, TOM 8, McGill's Ling Tea, Johns Hopkins's Semantics Lab, and everyone else I talked to along the way. Thanks also to anonymous reviewers for NELS 46, SALT 25 and AC 2015 for thought-provoking comments. Finally, I give special thanks to Luis Alonso-Ovalle and Bernhard Schwarz for extensive comments and advice, without which this work would not have been possible. All mistakes are my own. I gratefully acknowledge funding for this project from the FQRSC grant 2012 SE-144646, and the SSHRC grants 435-2012-0882, 435-2013-0592, 435-2014-1504, and 435-2013-0103, as well as support from McGill University.

©2017 Daniel Goodhue

This is an open-access article distributed under the terms of a Creative Commons Attribution License (<https://creativecommons.org/licenses/by/3.0/>).

1 Introduction

The primary concern of this paper is the felicity conditions on the use of epistemic modals. It has been argued by von Fintel & Gillies (2010) that *must* ϕ is felicitous only if the evidence for ϕ is indirect. To probe this claim, we will consider example pairs with the following profile: Two speakers, A and B, each say *It must be raining*. Suppose that A and B share identical perceptions, but nevertheless A's *must* ϕ utterance is intuitively felicitous while B's is not. As von Fintel & Gillies acknowledge, in order for their account to handle such example pairs, it must be claimed that what counts as direct evidence is context dependent and vague. While A and B share the same perceptual evidence E , only B takes it as direct evidence that it is raining. The question is, why? Under von Fintel & Gillies's indirectness account, no answer is forthcoming.

In this paper, I will argue that the indirectness account remains incomplete until a principled explanation of this context dependence is given (Section 2). In search of such an explanation, I will show that the most reasonable conclusion is that the felicity conditions of epistemic *must* ϕ utterances are not about directness of evidence at all, but are instead about knowledge: *Must* ϕ is felicitous only if ϕ is not known. I will refer to this generalization connecting *must* ϕ utterances and knowledge as *the epistemic account* (Section 3). I will argue that principled explanations of the difference between A and B can be offered under such an account. The conclusions I draw have ramifications for another debate about epistemic modals, namely whether or not *must* ϕ is strong. In alignment with Karttunen 1972, Kratzer 1991, Lassiter 2016, and Giannakidou & Mari 2016, I will argue that it is not.

Finally, I make overtures to two areas requiring future work in the wake of my arguments. First, I demonstrate that the epistemic account can be derived as a conversational implicature (Section 5). Second, I consider a remaining puzzle, the use of *must* ϕ in the conclusions of deductions, and demonstrate that the puzzle is a challenge for any account of epistemic modals. Then I suggest a way forward (Section 6).

We'll begin with some necessary background on epistemic modality.

1.1 The apparent weakness of epistemic *must*

Karttunen (1972) reports the intuition that (1a) does not entail (1b), and claims more generally that epistemic *must* ϕ does not entail ϕ .

Must ϕ is felicitous only if ϕ is not known

- (1) [From Karttunen 1972]
- a. John must have left. (*must* ϕ)
 - b. John has left. (ϕ)

If we thought that natural language epistemic *must* could be analyzed using the necessity modal from epistemic logic \Box , then the intuitions about the relationship between (1a) and (1b) are puzzling. This is because in classical epistemic logic a reflexive accessibility relation is generally used, thus $\Box\phi$ entails ϕ . So natural language *must* cannot be analyzed as \Box .

Of (1), Karttunen (p. 12) writes, “(1a) seems to say that the truth of *John has left* in some way logically follows from other facts the speaker knows *and some reasonable assumptions that he is willing to entertain*.” (italics mine) Kratzer (1991) proposes a semantics that implements these observations as follows: *Epi* is a contextually determined set of known propositions, thus $\bigcap Epi$ is the set of epistemically live possibilities, the epistemic modal base. $\bigcap Epi$ is restricted by a contextually determined set of reasonable assumptions (the stereotypical ordering source), which we’ll call *Norm*. The resulting set of possibilities are the most stereotypical ones, the ones that render the most propositions in *Norm* true. We’ll call this subset of $\bigcap Epi$ *Best*.¹ *Must* ϕ quantifies universally over *Best* and claims that ϕ holds throughout. Officially:²

$$(2) \quad \llbracket \text{must } \phi \rrbracket^{c, w} = 1 \Leftrightarrow \forall w' \in \text{Best} : \llbracket \phi \rrbracket^{c, w'} = 1$$

¹ Kratzer (1991) assumes that *Epi* and *Norm* (her $f(w)$ and $g(w)$ respectively) are provided by the context c . *Norm* induces an ordering on $\bigcap Epi$ as follows (being lower on the ordering means being more optimal with respect to the propositions in *Norm*):

$$(i) \quad \forall w, w' \in \bigcap Epi : w \leq_{\text{Norm}} w' \Leftrightarrow \{p \in \text{Norm} \mid w \in p\} \supseteq \{p \in \text{Norm} \mid w' \in p\}.$$

With this ordering in hand, *Best* is defined as follows:

$$(ii) \quad \text{Best}_{\bigcap Epi, \leq_{\text{Norm}}} = \{w \in \bigcap Epi \mid \forall w' \in \bigcap Epi [w' \leq_{\text{Norm}} w \rightarrow w \leq_{\text{Norm}} w']\}.$$

To determine $\text{Best}_{\bigcap Epi, \leq_{\text{Norm}}}$, $\bigcap Epi$ and \leq_{Norm} are required, which in turn are only made available thanks to the functions f and g provided by c , and a world of evaluation w to act as input to those functions. So $\text{Best}_{\bigcap Epi, \leq_{\text{Norm}}}$ is dependent on c and w for its content. However, I will leave these subscripts off and just refer to *Best* in the body of the text.

² Kratzer’s (1991) denotation for *must* ϕ (p. 644) is more complex than that in (2) because she does not make the limit assumption. Without this assumption, there is no set of *Best* worlds. See von Stechow & Heim 1997–2023: p. 61–62 for some discussion and references.

Since the actual world w may be non-stereotypical (i.e., not in *Best*), *must* ϕ can be true even when ϕ is false. Thus Kratzer explains the intuitions about (1), solving Karttunen's problem.

1.2 The apparent indirectness of epistemic *must*

von Fintel & Gillies (2010) offer a solution to Karttunen's problem from a different angle. They want to maintain that *must*, like \Box , is a universal quantifier ranging over all epistemic possibilities in $\cap Epi$, and thus that *must* ϕ entails ϕ , despite Karttunen's reported intuitions about (1). To do this, they focus on a different empirical observation, based on the following two utterances of *must* ϕ .

- (3) [Adapted from von Fintel & Gillies 2010: p. 353]
 Billy is looking out the window at the pouring rain.
 a. #It must be raining.
- (4) [Adapted from von Fintel & Gillies 2010: p. 353]
 Billy sees someone enter the building holding a wet umbrella, but she cannot see outside herself.
 a. It must be raining.

von Fintel & Gillies (vF&G) argue that the reason that (3a) is infelicitous while (4a) is felicitous is that *must* ϕ has an indirectness requirement: *Must* ϕ is felicitous only if the evidence for ϕ is indirect. They have been followed in this view on indirectness by several researchers since (e.g., Ozturk & Papafragou 2015, Matthewson 2015, Lassiter 2016, Mandelkern 2016), though there has been disagreement over whether *must* ϕ is strong.

vF&G say that epistemic necessity modals in all languages they have checked convey this indirectness signal, thus they would prefer to derive it as a quantity implicature (Grice 1989, Geurts 2010, Franke 2011, Frank & Goodman 2012). However, they argue that it is not clear how to do this since there is no stronger expression $O\phi$ that entails *must* ϕ and that requires the evidence for ϕ to be direct. Thus they implement the indirectness requirement as a presupposition as follows: *Direct* is a special subset of *Epi* containing only propositions that are directly known. ϕ cannot be entailed by any single proposition in *Direct*. If this presupposition is met, then *must*

Must ϕ is felicitous only if ϕ is not known

quantifies universally over $\bigcap Direct$ and says that ϕ holds throughout. Officially:³

- (5) a. $\llbracket \text{must } \phi \rrbracket^{c,w}$ is defined only if $\forall \psi \in Direct: \psi \not\subseteq \lambda w. \llbracket \phi \rrbracket^{c,w}$
 b. If defined, $\llbracket \text{must } \phi \rrbracket^{c,w} = 1 \Leftrightarrow \forall w' \in \bigcap Direct: \llbracket \phi \rrbracket^{w'} = 1$

Since *Direct* is a subset of *Epi*, every proposition in *Direct* is true in w , thus $\bigcap Direct$ contains w , and ϕ is predicted to hold in w . Thus, vF&G's semantics predict that *must* ϕ entails ϕ .

If *must* ϕ feels weaker than ϕ , they say, it's not because it is actually weaker, it's just because the evidence for ϕ is indirect, which gives rise to a non-confidence intuition. Thus they claim to explain the intuitions about (1) as indirectness or non-confidence intuitions while maintaining that *must* ϕ is strong.

1.3 The apparent ignorance of epistemic *must*

While both Kratzer 1991 and von Stechow & Gillies 2010 offer explanations of the intuitions surrounding *must* ϕ in (1), only vF&G's indirectness account offers an explanation of the contrast in felicity intuitions in (3) and (4). Giannakidou & Mari (2016: G&M) modify Kratzer's account to offer a different explanation of these felicity intuitions. G&M claim that the felicity conditions of epistemic modals are not about directness, but are instead a requirement that the speaker does not know the prejacent. Thus they add a presupposition to Kratzer's semantics in (2) that requires that $\bigcap Epi$ does not entail ϕ . Officially:

- (6) a. $\llbracket \text{must } \phi \rrbracket^{c,w}$ is defined only if $\bigcap Epi \not\subseteq \lambda w. \llbracket \phi \rrbracket^{c,w}$
 b. If defined, $\llbracket \text{must } \phi \rrbracket^{c,w} = 1 \Leftrightarrow \forall w' \in Best: \llbracket \phi \rrbracket^{w'} = 1$

G&M's account inherits Kratzer's explanation of (1) as weakness. (3a) is predicted to be infelicitous because Billy knows for sure that it is raining, while (4a) is predicted to be felicitous because Billy does not know ϕ for sure. Thus, just as vF&G recharacterize the intuitions about (1) from weakness to

³ vF&G's presupposition includes a second conjunct requiring that no proposition in *Direct* contradicts the prejacent either, which is relevant to modals embedded under negation (e.g., *There don't have to be two reds*). These concerns aren't relevant here, so I leave it out to simplify presentation.

non-confidence, G&M recharacterize the intuitions about (4). For vF&G, Billy indirectly knows ϕ for sure in (4), while for G&M, Billy does not know ϕ .

We are left with two distinct theories that capture the empirical observations discussed so far. To argue in favor of their account, G&M introduce the following examples.

(7) [From G&M 2016]

You and your sister were out of touch for a couple of years. Today she calls you on the phone to catch up. She tells you that her daughter Maria plays the piano. Later, you tell your husband:

a. #Maria must play the piano.

G&M claim that in (7), the evidence for the prejacent *that Maria plays the piano* is indirect since the speaker is told about it. Therefore, they say that the indirectness account does not explain the infelicity intuition. However, if we consider carefully how vF&G's account works, they say that trustworthy reports count as direct evidence. Therefore, the prejacent of (7a) is included in the *Direct* set, so (7a) is correctly predicted to be infelicitous.

Consider (8):

(8) [From G&M 2016]

You and your sister were out of touch for a couple of years. Today you visit her for the first time. As she shows you around her apartment, you see that there is a piano. Later, you tell your husband:

a. Maria must play the piano.

G&M claim that in (8), the evidence for the prejacent is direct because the speaker sees the piano, therefore the indirectness account incorrectly predicts (8a) to be infelicitous. In (8), G&M are grappling with the issue of what counts as direct evidence under the indirectness account (an issue we will be grappling with ourselves in a moment). Does seeing a piano in someone's house count as direct evidence that they play the piano? I suspect that vF&G would argue for a tighter connection between direct observations and direct evidence in the *Direct* set. Just as seeing a wet umbrella does not produce direct evidence for the proposition *that it is raining*, seeing a piano at Maria's does not produce direct evidence for the proposition *that Maria plays the piano*. Instead, these observations merely produce direct evidence for the propositions *that there is a wet umbrella* and *that Maria has a piano*. There-

Must ϕ is felicitous only if ϕ is not known

fore, I believe that (8a) is correctly predicted to be felicitous by the indirectness account.

Nevertheless, G&M's account explains (7) and (8) just as well as the indirectness account. They argue that (7a) is infelicitous because the speaker knows the prejacent, while (8a) is felicitous because the speaker does not know the prejacent. So far then, we are still left with two accounts of the felicity conditions of epistemic modals and no evidence that clearly distinguishes between them.⁴

In the next section, I will introduce new evidence to distinguish between these accounts.

2 Identical perceptions, differing felicity intuitions

von Fintel & Gillies's (2010) indirectness account depends on the contents of *Direct* to make predictions about felicity intuitions. Thus, we need to know how propositions enter *Direct*. According to vF&G, *Direct* contains propositions arising from direct observation.⁵ But how does this work exactly? To be sure, exploring the intimate relationship between perceptions and evidence

⁴ G&M provide some data from Greek and Italian, for example (i), that is incompatible with the strength of vF&G's account (G&M argue that future marking in Greek and Italian conveys epistemic necessity). Lassiter (2016) provides equivalent English corpus data, for example (ii).

- (i) I Ariadne tha ine arosti, ala dhen ime ke endelos sigouri
the Ariadne FUT is sick, but not be.1sg and absolutely sure
'Ariadne must be sick, but I am not entirely sure.' Greek
- (ii) Must be an old DTS diesel setup but I'm not certain.

While such examples are clearly a challenge for vF&G's claim that *must* ϕ entails that ϕ is known for sure, the strength component and the indirectness presupposition of vF&G's account are independent in the sense that the indirectness presupposition could be combined with a weak truth conditions for *must* ϕ . This would predict (i) and (ii) to be felicitous. In other words, these examples do not demonstrate that the indirectness presupposition makes incorrect predictions. Nor do they demonstrate that *must* ϕ requires ϕ not to be known (G&M's position), just that it is compatible with ϕ not being known (Kratzer's/Lassiter's position). Lassiter (2016) discusses the relationship between weakness, knowledge and certainty in more depth. Note that compatibility with lack of knowledge contradicts one aspect of vF&G's position, namely that *must* ϕ requires ϕ to be known (strength). However, there may be ways for a defender of a strong *must* to explain such examples, such as claiming that the modal base shifts in the second conjuncts of (i) and (ii).

⁵ vF&G also say that *Direct* contains trustworthy reports, and Matthewson (2015) points out that *Direct* also needs to contain what she calls general reasoning conditionals (e.g., (4a)

is not typical fodder for semantic research. Nevertheless, this is where the indirectness account leads us. If we want to know what predictions this theory makes about epistemic *must* statements, then we need to consider perceptions and evidence.

The most straightforward answer is that there is a one to one mapping between perceptions (direct observations) and propositions in *Direct*. Let's probe this idea, *the one to one theory*, with the following pair of examples in which each of the speakers has exactly the same perceptual experience.

- (9) Phil is cooking chicken and peas for his family. When the timer goes off, he checks the chicken's temperature and discovers it is done. He tastes the peas and they are also ready. The table is already set. Phil's daughter comes in and says, "Is dinner ready?" Phil says:
- a. #Dinner must be ready.
- (10) Phil is cooking dinner for his family and his friend Meryl. He had to step out in a hurry, and instructed Meryl as he left: "Please turn the peas off when they are done, and take the chicken out of the oven when the temperature is right." When the peas are done, Meryl turns the burner off, and when the chicken is done, she removes it from the oven. She has also seen that the table is set. She wonders whether Phil was planning to make anything else, for example a salad, but Phil didn't mention anything. Phil's daughter comes in and says, "Is dinner ready?" Meryl says:
- a. Dinner must be ready.

(9a) is intuitively infelicitous while (10a) is intuitively felicitous. For the indirectness account to explain the asymmetry, it would need to be claimed that Phil has direct evidence for the prejacent ϕ while Meryl does not. However Phil and Meryl share identical perceptions, so according to the one to one theory, their *Direct* sets should contain exactly the same propositions. Thus, the one to one theory cannot be right.⁶

requires *Direct* to contain *if there is a wet umbrella, then it is raining* in order to be true). We'll leave these issues aside since neither trustworthy reports nor general reasoning conditionals are crucial to any of the felicity intuitions we are about to consider.

⁶ As mentioned in footnote 5, we are restricting our attention to propositions derived from perceptual experience. Phil's and Meryl's *Direct* sets may differ in terms of propositions corresponding to trustworthy reports or general reasoning conditionals. In fact, Phil clearly (directly) knows a bit of crucial information that Meryl does not, namely, he knows what is for dinner. This will be discussed in Section 3. The point here is that neither the proposition

Must ϕ is felicitous only if ϕ is not known

For the indirectness account to be maintained, it will have to be claimed that two agents with identical perceptions can nevertheless have different direct evidence in the form of different contents in *Direct*. In order to avoid stipulating that *Direct* has the right contents to predict our felicity intuitions, an explanation of how perceptions become propositions in *Direct* needs to be given.

vF&G acknowledge the complicated mapping between perceptions and direct evidence, claiming that what counts as “direct” is vague and context dependent (p. 369-370). The evidence that drives them to make this claim can be seen by comparing Billy in (3) with the skeptical epistemologist in (11):

- (3) [Adapted from von Fintel & Gillies 2010: p. 353]
Billy is looking out the window at the pouring rain.
a. #It must be raining.
- (11) [Adapted from von Fintel & Gillies 2010: p. 370]
A professional epistemologist, while on vacation in Seattle, looks out the window at the pouring rain. She says:
a. It must be raining.

If we bear in mind that professional epistemologists are apt to be skeptical of their own perceptions due to their training (at least in certain contexts), then (11a) is intuitively felicitous. The differing intuitions about the *must* ϕ utterances of Billy and the epistemologist is yet another way of demonstrating the point already made by (9) and (10) above: The contents of *Direct* cannot be tied to perceptual experience too tightly.

I will discuss how these pairs of examples differ from one another in more detail below, but first note that we don’t need to restrict ourselves to the relatively exotic case of skeptical epistemologists. An example with a layperson:

- (12) Hillary is in her office and sees falling rain out the window. She received an e-mail that morning saying that a Hollywood movie would be filmed outside that day, and that if it didn’t rain they would be

corresponding to this information, nor any proposition corresponding to Phil’s trustworthy reports or general reasoning conditionals entails the prejacent *that dinner is ready*, as would be required by (5a) if vF&G’s indirectness account were to make the right prediction. So we can see that Phil’s *Direct* needs to somehow contain a proposition that Meryl’s does not, despite their identical perceptions and regardless of their trustworthy reports or general reasoning conditionals.

making fake rain, though the filming isn't supposed to start until 5 pm. Hillary looks at the clock, which reads 4:52 pm.

- a. It must be raining.

Just like the epistemologist, and unlike Billy, Hillary's *must* ϕ utterance is intuitively felicitous. All three speakers see the same exact thing, falling rain out the window. Thus, in order to accept the indirectness account, we need to accept that what counts as direct evidence is context dependent.

Clearly an explanation for the proposed context dependence of *Direct* is needed. Otherwise, the theory does not make predictions about the felicity of *must* ϕ utterances. But no explanation is offered in von Fintel & Gillies 2010. In its present form, the indirectness account simply stipulates that speakers whose *must* ϕ utterances are intuitively infelicitous have direct evidence, while those whose *must* ϕ utterances are intuitively felicitous have indirect evidence.⁷ The stipulations seem intuitive for those example pairs in which the agents have different perceptions, but not for pairs where the agents have identical perceptions. Perhaps there is a satisfying explanation for why a single perception can lead to *Direct* evidence for one agent and not another, but at the moment we do not have such an explanation.

In search of an explanation, we turn to the epistemic account.

3 The epistemic account

Intuitively, Phil in (9) knows for sure that dinner is ready, while Meryl in (10) does not. This difference in knowledge depends on another piece of knowledge that Phil has but Meryl lacks, namely Phil knows what is for dinner. Likewise, the difference between Billy in (3) and the skeptical epistemologist in (11) has to do with knowledge. This is what makes skeptical epistemologists interesting, they lack the kind of everyday knowledge that laypeople have (discussed at length in Lewis 1996, to be discussed below). Even rain perceptions aren't quite good enough to allow them to conclude that it is

⁷ In fact, vF&G (p. 370-371) say that they do not want to simply stipulate indirectness by labeling some propositions as INDIRECT. "Instead, we want the evidential signal to emerge from combining the basic ingredients of the semantics of modals." They achieve this by defining an indirectly known proposition as one that is not entailed by any single proposition in *Direct*, but that is entailed by \bigcap *Direct*. While this definition of indirectness is a cut above labeling a proposition "INDIRECT", it relies on another stipulation, namely labeling some propositions as DIRECT. Thus we are still left with an account that, at bottom, stipulates in/directness.

Must ϕ is felicitous only if ϕ is not known

raining. Hillary in (12) also has reason to doubt her perceptions, thus also fails to know the prejacent for sure.

It seems clear that there is a generalization that holds between the felicity conditions on *must* ϕ utterances and knowledge of ϕ :

- (13) *The epistemic account:*
Must ϕ is felicitous only if ϕ is not known.

Giannakidou & Mari's (2016) account is of course one way of spelling this generalization out as a presupposition. The generalization could also be captured as a conversational implicature, which I will explore in Section 5. Different ways of implementing the generalization may make predictions that differ in interesting ways. However, the main insight of this paper is not a particular implementation of (13), but (13) itself. The felicity conditions of epistemic modals are about whether the prejacent is known.

I ended the last section saying that it may be possible to explain the empirical contrasts via indirectness, but that some deeper explanation of what it means to be direct is needed. Could knowledge just be the terms in which in/directness is defined, that is, *direct* =_{def} *known* and *indirect* =_{def} *not-known*? If we are considering von Stechow & Gillies's (2010) particular indirectness account, then the answer is no. The epistemic account is incompatible with vF&G's claim that *must* ϕ is strong, a universal quantifier over what is known. The two lead to a contradiction. Thus, we can see that adopting the epistemic account leads fairly directly to taking *must* ϕ to be weak.⁸

Lassiter (2016) adopts vF&G's indirectness requirement, but unlike them, he combines it with a weak *must*. Under this state of affairs, could we take

⁸ There is a veridical theory of *must* ϕ that is compatible with the epistemic account, although the result is somewhat strange. This particular veridical theory, considered in Lassiter 2016, combines a weak semantics for *must* ("weak" in that *must* ϕ merely requires ϕ to have a high likelihood given the evidence) with a conjunct requiring ϕ to be true. As long as asserting some proposition does not require knowing that proposition, then one could assert this weak but veridical *must* ϕ without knowing ϕ , and *must* ϕ would entail ϕ . Thus the epistemic account and veridicality could be simultaneously maintained using this semantics. This approach would lead to the odd claim that a speaker who asserts *must* ϕ never knows for sure whether their utterance is true, since they are required by the epistemic account not to know for sure whether ϕ is true. This suggests that we wouldn't want to combine the epistemic account with a veridical theory of *must* ϕ . I will return to these issues below in footnote 12 when proposing an implicature derivation of the epistemic account. This will lead us to adopt the knowledge norm of assertion and conclude that *must* ϕ cannot be veridical.

knowledge to be the terms in which we understand indirectness? Maybe, but now we are just discussing terminology. If the more fundamental concept in terms of which we understand indirectness is knowledge, then we might as well do away with any discussion of indirectness, and just call the phenomenon what it is. We don't gain anything by continuing to use the language of in/directness if the underlying concept that explains the intuitions is knowledge.

The advantage of the epistemic account is that we have intuitions about the knowledge of agents in the crucial example pairs, and these intuitions systematically correspond to intuitions about the felicity of *must* ϕ utterances, as spelled out in (13). This systematic correspondence makes predictions. The indirectness account, on the other hand, does not make clear predictions because it does not explain why some agents count their evidence as direct while others do not, and we do not have intuitions about whether an agent has direct evidence or not in the crucial examples, or at least nowhere near as clear-cut intuitions as we have about knowledge.

In Sections 3.1 and 3.2, I will demonstrate how the context conspires to render ϕ known for some agents and not others in our context pairs. This discussion is guided by the theory of knowledge developed in Lewis 1996, though we will not explicitly consider his work until Section 3.2. Lewis considers example pairs like those already discussed, however he does not consider *must* ϕ utterances. Instead, his aim is to explain our intuitions about knowledge. Lewis helps us to see that some agents fail to know because they know too little, while others fail to know because they know too much. This will help us come to grips with the role that evidence plays in knowing ϕ , which in turn helps us understand the use of *must* ϕ utterances.

3.1 Knowing too little vs. eliminating $\neg\phi$ possibilities

Take Phil from (9) and Meryl from (10). Phil has access to some extra information in that he knows what is for dinner: chicken and peas and nothing else. Thus he knows that if both of those are done and the table is set then dinner is ready. When this conditional is conjoined with the knowledge arising from Phil's perceptions, Phil is able to eliminate all possibilities in which dinner is not ready. Thus Phil knows that dinner is ready, and by (13), he cannot felicitously say (9a). Putting this in the language of the possible worlds semantics for modals discussed in Section 1, when Phil's knowledge of what is

Must ϕ is felicitous only if ϕ is not known

for dinner is conjoined with the other propositions in *Epi* derived from Phil's perceptions, the resulting $\bigcap Epi$ entails the prejacent, *that dinner is ready*.

Meryl, on the other hand, lacks this extra knowledge, thus she knows too little. If Phil had told her what was for dinner, then she would have had the proper information via trustworthy report. Without it, she knows that the chicken and peas are done and that the table is set thanks to her perceptions, but that information together does not entail the prejacent, thus she doesn't know for sure whether dinner is ready. According to (13), she can felicitously say (10a).

While Meryl does not know ϕ , she seems to be almost certain that ϕ as demonstrated by her claim that *must* ϕ , which means that ϕ follows from her reasonable assumptions and other things she knows, according to Karttunen's (1972) gloss. In order for this to work, she has to be assuming something like, if someone gives you some instructions for making dinner before rushing off, then the instructions are complete. Meryl doesn't know for sure that there is no other dish, like a salad, but if she assumes the instructions are complete, then she can conclude that dinner is ready. She could even mention this assumption explicitly by saying, "Phil probably told me everything I had to do, so dinner must be ready." Thus it is clear why Meryl's *must* ϕ utterance gives rise to a weakness intuition: She is relying (potentially explicitly) on a reasonable assumption that could turn out to be false.

Unlike the indirectness account, according to the epistemic account speakers who can felicitously say *must* ϕ and those who cannot both can combine information to get to ϕ . The difference between them is not about directness at all, but about whether ϕ known, which in this case boils down to whether ϕ follows only from known facts, or whether the conclusion that ϕ also depends on defeasible assumptions.

3.2 Knowing too much vs. ignoring $\neg\phi$ possibilities

Consider Billy in (3), the skeptical epistemologist in (11), and Hillary in (12), who all see falling rain. In some sense, the skeptical epistemologist and Hillary each have some extra information that Billy does not have, but, unlike Phil, this extra information causes them not to know the prejacent ϕ . To see how this works, let's consider Lewis's (1996) context dependent theory of knowledge, which is meant to explain the difference between skeptical epistemologists and laypeople.

- (14) The speaker knows that $\phi \Leftrightarrow$ The speaker's evidence eliminates every possibility in which $\neg\phi$ — except for those possibilities that we are properly ignoring.

We have supposed our three speakers all have identical rain perceptions. Do they have identical evidence arising from those perceptions? According to Lewis, yes. Their identical perceptions give rise to identical evidence E with identical propositional content ϕ . An agent's evidence E eliminates every possible world in which the agent has evidence E with content different from ϕ . Crucially, this is not the same as claiming that E eliminates every possibility in which $\neg\phi$. There are still possibilities in which it is not raining that E has not eliminated. For instance, there are farfetched possibilities in which the agent is suffering from a delusion or being tricked into having rain perceptions when it is nevertheless not raining. These farfetched possibilities cannot be eliminated by E , they are still live possibilities in which the agent is having rain perceptions despite that it is not actually raining. That is, in these farfetched possibilities, the agent has perceptual evidence E with propositional content ϕ despite that $\neg\phi$ holds in each of these farfetched possibilities.

If Billy, the epistemologist and Hillary all have identical evidence E eliminating exactly the same possibilities, then how do they differ? According to Lewis, Billy properly ignores the farfetched possibilities while the epistemologist and Hillary do not. After removing all of the possibilities that Billy eliminates or ignores, the only live possibilities left are those in which it is raining. Thus Billy knows that it is raining.

The epistemologist and Hillary make the same eliminations via the same perceptual evidence, but they are left with some possibilities in which it is not raining that they cannot properly ignore. Why can't they ignore them? For the epistemologist, it is because her training as an epistemologist gets in the way. She is consciously aware of the farfetched possibilities, and to be aware of them is to not ignore them. So, she doesn't quite know for sure that it is raining. Billy knows ϕ because she ignores farfetched possibilities in which $\neg\phi$. The epistemic account predicts her *must* ϕ utterance to be infelicitous. But the epistemologist knows too much to know ϕ . She cannot ignore these possibilities as a result of being aware of them. Thus her *must* ϕ utterance is felicitous.⁹

⁹ There is a connection between the ignorance that enables knowledge in Lewis's account, and unawareness as discussed in [Franke & de Jager 2011](#). In one of their primary examples, Bo

Must ϕ is felicitous only if ϕ is not known

We are now able to see that Hillary in (12) is roughly like the skeptical epistemologist. She can be sure that she has rain perceptions, but she cannot take the usual step of concluding *that it is raining* from her perceptual evidence. This is because some farfetched possibilities in which her perceptions are an illusion are unignorable due to the e-mail she has received. That is, even though a layperson with rain perceptions would normally ignore the possibility that it was an illusion, this behavior quickly disappears in a context in which illusions are highly likely. Thus Hillary does not know ϕ , and by (13) can felicitously say *must* ϕ .

Despite her lack of knowledge, Hillary seems fairly sure it is raining. This is because she can combine her rain perceptions with other things she knows (the time, when the fake rain is supposed to start), and she can make the reasonable assumption that if people say they will do something at a certain time, then they don't start until that time. This assumption is obviously not always true, for example the film crew could be testing their equipment out early. But stereotypically it is. Thus Hillary infers ϕ from facts and reasonable assumptions, as indicated by her *must* ϕ utterance.

We can see now the role that evidence plays in knowing ϕ , and how it relates to the use of *must* ϕ utterances. An agent's evidence eliminates possibilities in which the agent does not have that evidence. This evidence needs to be combined with some ignoring in order to know ϕ . Lewis's theory allows us to take the reasonable position that agents with identical perceptions have

is missing her keys. She is unaware of (ignores) the possibility that the keys might be in her car. But she knows they could be in one of several places in her house. Thanks to her unawareness, she feels certain that they are in the house. Thus, after some looking reveals that the keys are not in any of the possible house-locations, frustration results. There is an interesting difference between this case and Billy's owing to differences in the relevance or likelihood of the ignored possibilities. If someone suggests to Bo that her keys might be in the car, she adjusts her beliefs and goes to check the car. If someone suggests to Billy that her perceptions may be non-veridical or that the rain may be fake, she is apt to brush this suggestion off and continue to know that it is raining. Some kinds of unawareness or ignoring may be safer than others. Indeed, Lewis claims that ignoring is a weak way to know, which can be defeated by merely attending to ignored possibilities, as demonstrated by Bo and her missing keys. However Blome-Tillmann (2009: p. 246-249), who builds on Lewis 1996, discusses more stable ignoring like Billy's. He argues that the contrast is not ignoring vs. attending to (being aware that), but ignoring vs. taking seriously, which he cashes out as pragmatic presupposition (p. 253). It seems clear to me that one can ignore possibilities either through inattention/unawareness or through lack of credence. Blome-Tillmann's use of pragmatic presupposition may capture both of these, however, as Franke & de Jager (2011) suggest, more work may be needed to understand the connection between awareness and presupposition.

identical evidence. We do not need to claim that Billy has direct evidence for ϕ while the skeptical epistemologist and Hillary have indirect evidence. They all have the same evidence owing to their identical perceptions, and only Billy knows ϕ because only she can ignore farfetched possibilities in which she has evidence for ϕ but $\neg\phi$ holds. This story explains our intuitions about what these agents know, and it also shows us that the use of *must* ϕ does not depend on a distinction between direct and indirect evidence, which has no status in the theory. Rather it depends on what is known.

3.3 A note on logical omniscience

The epistemic account says that *must* ϕ is felicitous only if ϕ is not known. We have just seen two ways that ϕ can fail to be known, knowing too little and knowing too much. As predicted by the epistemic account, the felicity intuitions about *must* ϕ utterances track knowledge in each of these cases. Are there other ways that ϕ can fail to be known? If so, will the intuitions about *must* ϕ utterances track these failures in knowledge as predicted?

Consider the logical omniscience problem, which might be viewed as a special subcase of knowing too little, but which nevertheless has its own unique signature. The logical omniscience problem arises when ϕ follows from what is known, but the agent does not recognize this consequence. For example, perhaps the inference to ϕ is rather complex, or even though two known propositions together entail ϕ , the speaker hasn't thought about combining them yet. The question is, if a speaker fails to know ϕ due to this problem, will her *must* ϕ utterance be felicitous as predicted by (13)?

To test this prediction we need an example where the prejacent is entailed by the speaker's knowledge, the speaker fails to see this entailment, and nevertheless the speaker has a high degree of confidence that the prejacent is true because it follows from a combination of other things she knows and her reasonable assumptions.

Consider the following examples, adapted from examples suggested by Michael Franke (p.c.).

- (15) Profs. SmallClass and BigClass are grading final exams. Prof. SmallClass has five students total and has so far graded three exams, while prof. BigClass has fifty students total and has so far graded thirty exams. The professors have the following property in common: For any student in her class, if asked whether she has graded that student's

Must ϕ is felicitous only if ϕ is not known

exam, she can confidently give the correct answer. But if you asked each of them to write a list from memory of which students they have graded so far, only Prof. SmallClass could do it. Prof. BigClass would fall a bit short. Each of their spouses calls them and asks how much more work they have to do. They reply as follows:

- a. Prof. SmallClass: # I must have graded more than half (by now).
- b. Prof. BigClass: I must have graded more than half (by now).

While both professors can demonstrate knowledge of whose exams they have graded, only prof. SmallClass seems to know how many exams total she has graded so far. Prof. BigClass lacks this knowledge, and as such she does not know the prejacent. The epistemic account in (13) predicts (15a) to be infelicitous, and (15b) to be felicitous, and these predictions are borne out. The logical omniscience problem is another example demonstrating that when knowledge of the prejacent fails, epistemic *must* is felicitous, in line with the epistemic account.

As is well known (e.g., Fagin & Halpern 1987, Fagin et al. 1995, Yalcin 2018), the logical omniscience problem is a problem in that possible worlds models of epistemic reasoning typically predict agents to immediately know all deductive consequences of their knowledge, but real humans do not display this behavior. For example, consider the outcome for (15) if we apply the possible worlds implementation of the epistemic account proposed by Giannakidou & Mari (2016) in (6). Given the professors' ability to correctly say whose exams they have graded, it seems reasonable to say that for each student, they know whether they have graded that student's exam. This information, when combined with the knowledge of their total class sizes, produces an $\bigcap Epi$ for each professor that entails the prejacent *I have graded more than half*. Thus, while the epistemic account in (13) makes the right predictions, the implementation of it in (6) makes the right prediction for (15a), but not for (15b).¹⁰

¹⁰ The problem is that a set of propositions like *Epi* and its conjunction $\bigcap Epi$ do not model the fact that humans sometimes fail to see the consequences of what they know. Taking *Epi*'s conjunction is like scanning through a spreadsheet and manipulating data to see all of its consequences. People can do this just fine if the amount of data is small. But we struggle if there is a lot of data. Extra-mental tools (like spreadsheets) are needed. The model in (6) does not make this distinction.

The most straightforward way to fix the particular implementation in (6a) is to apply the idea of compartmentalized belief (Stalnaker 1984, discussed also in Lewis 1996 and Yalcin 2018): Agents only recognize the logical consequences of certain members of the powerset

The literature on the epistemological puzzle of the logical omniscience problem makes clear that a better account of how people reason with knowledge is needed for reasons independent of epistemic modality. Whatever the best solution is, it will need to be imported into the semantics of epistemic modals. The key point here is that the generalization in (13) between the intuitions about knowledge and the felicity conditions on *must* ϕ utterances, which is the main thesis of this paper, is further supported by this class of examples.

4 In the wake of the epistemic account

So far, I have introduced the epistemic account and argued that it offers a preferable explanation of the felicity conditions of epistemic modals.

In the remainder of the paper, I will discuss two issues that arise in the wake of the epistemic account. First, in Section 5, I will ask whether we can explain why there is a generalization holding between knowledge and *must* ϕ utterances. I suggest that the epistemic account is amenable to a derivation via quantity implicature, fulfilling an expectation expressed in von Fintel & Gillies 2010.

Second, in Section 6, I will discuss a serious challenge to the epistemic account, the fact that *must* ϕ utterances can appear in the conclusions of deductions. I will show that this is in fact a challenge to the indirectness

of *Epi*, with availability of a set determined in this case by how large it is. *Must* ϕ is felicitous only if ϕ is not entailed by one of the available subsets of *Epi*. Prof. SmallClass has access to the consequences of the relevant subset in *Epi* while prof. BigClass does not.

However it is already known that such a solution won't solve the problem of logical omniscience more generally. For example, what if, for a particular subset of *Epi*, an agent recognizes some of its logical consequences and not others? Moreover, as Yalcin (2018) points out, the problem of logical omniscience is actually a cluster of problems and this solution does not solve all of them. For example, an agent may fail to recognize that one proposition they know entails another. In this guise, the logical omniscience problem also poses a challenge to vF&G's implementation of the indirectness account in (5a) (i.e., a directly known proposition ψ entails ϕ , thus *must* ϕ is predicted to be infelicitous). But if the agent fails to recognize the entailment, then the *must* ϕ utterance may be intuitively felicitous. For example, an ancient who knows directly that the celestial body they are looking at is the morning star, and who felicitously says (because they are just now in the process of working out a hunch that the morning star and the evening star are the same body) *That must be the evening star*. The utterance is intuitively felicitous, though it is predicted to be infelicitous by (5a).

Must ϕ is felicitous only if ϕ is not known

account as well, and I will suggest a way of dealing with the problem that solves it for everyone. In [Section 7](#), I will offer some concluding remarks.

5 An implicature account is now possible

[von Fintel & Gillies \(2010\)](#) say that the indirectness presupposition is a placeholder for an eventual explanation of why the proposed evidential signal of epistemic modals is persistent cross-linguistically. They write (p. 367), “...one would suspect and hope that the evidential signal can be derived as a conversational implicature that is non-detachable in Gricean terms.” As already mentioned above, such a derivation would depend on the existence of a stronger competitor $O\phi$ that is only appropriate with directly known prejacent. Here is a sketch of how this derivation would work:

(16) [Adapted from [vF&G \(2010: p. 367\)](#)]

A sketch of the desired implicature account:

1. *Must* is a universal quantifier over what is known.
2. *Must* competes with a stronger expression O that is only appropriate with directly known prejacent.
3. Choosing *must* ϕ instead of $O\phi$ implicates (via standard quantity implicature reasoning) that ϕ is not known directly.

The problem with trying to derive indirectness as a quantity implicature is that there is no stronger competitor $O\phi$ with the right properties. Thus [vF&G](#) fall back on a presuppositional implementation that they view as a placeholder for an eventual conversational analysis.¹¹

¹¹ [Mandelkern \(2016\)](#) offers a derivation of the felicity conditions of epistemic modals while maintaining the indirectness account by seeking to explain indirectness in terms of redundancy. He claims that *must* ϕ utterances require a salient shared argument in support of ϕ (for independent reasons). If that argument’s support of ϕ is so mutually obvious to the interlocutors that it renders an utterance of ϕ redundant, then the utterance is infelicitous. Thus the ban on mutual obviousness enforces indirectness. However, as [Mandelkern \(p. 16\)](#) notes, this approach leaves open the question of “what counts as a ‘mutually obvious’ inference”. So this account swaps one poorly understood concept, in/directness, for another, non/redundancy. Whatever label one chooses to use, an explanation for the contrasting intuitions in [Section 2](#) is still owed. And just as with the indirectness account, no explanation is forthcoming under this account.

See also [Faller 2012: p. 310](#) for another take on how to derive the felicity conditions of English epistemic modals that relies on a new notion of “evidential strength”.

However, if the epistemic account I have argued for is right, then the requirements for $O\phi$ are not what vF&G thought, O does not need to convey that ϕ is known directly. Moreover, the goal of an implicature account is no longer to explain a signal of evidential indirectness. The goal now is to explain an inference about the epistemic modal base, that ϕ is not known. In light of these changes, we can imagine a stronger alternative $O\phi$ that makes an implicature account possible, at least in principle. Here is how it would work:

- (17)
1. *Must* is a universal quantifier over what is known as restricted by reasonable assumptions (*Best*).
 2. *Must* competes with a stronger expression O that quantifies universally over what is known ($\cap Epi$).
 3. Choosing *must* ϕ instead of $O\phi$ implicates (via standard quantity implicature reasoning) that ϕ is not known.

If we can find a suitable competitor $O\phi$, then the sketch in (17) should work just fine. If *must* ϕ is used instead of $O\phi$, a hearer could reason that the speaker did not use the more informative $O\phi$ because it is not true, that is, because ϕ is not known. When ϕ is known, the hearer expects the speaker to use $O\phi$, not *must* ϕ (Grice 1989).

I believe that an assertion of the bare prejacent ϕ is a promising candidate for $O\phi$. First, it is a valid alternative to *must* ϕ structurally speaking. Katzir (2007) argues that alternatives to ψ are those structures that can be produced from ψ via deletions, contractions and replacements. *Must* ϕ can be transformed into ϕ via deletion, therefore, ϕ is a valid alternative to *must* ϕ .

Of course, the truth of ϕ does not depend on the contents of $\cap Epi$ or *Best*, so it does not, by itself, entail *must* ϕ . However, it has been argued that the norms of assertion are such that one can only assert what one knows (see Williamson 1996 for a defense, and Matt Weiner 2007 for a critical overview). If the knowledge norm is correct, then an assertion of ϕ entails *must* ϕ for that speaker, thus it would suffice in the role of $O\phi$ in (17). A listener can reason that if the speaker chose to say the weaker *must* ϕ , it is because what is known does not entail ϕ , or in terms of (6), $\cap Epi$ does not entail ϕ .¹²

¹² We can see now that deriving the epistemic account as a quantity implicature necessitates abandoning Lassiter's (2016) proposed veridical but weak *must*, discussed in footnote 8 above. This is because veridicality combined with the knowledge norm and the epistemic

Must ϕ is felicitous only if ϕ is not known

Whether the knowledge norm of assertion is correct is still a hotly debated open question. Thus whether the implicature derivation proposed here is correct in all its details must also remain an open question. However, I believe that the arguments in favor of the epistemic account put forth in this paper, combined with the desirability of deriving the epistemic account as a conversational implicature, and the fact that the knowledge norm of assertion is required for the quantity implicature derivation to go through, together constitute an argument in favor of the knowledge norm of assertion. That is, the fact that assertions of ϕ feel stronger than assertions of *must* ϕ , and that quantity implicature reasoning elegantly explains this intuition as well as those about the examples from [Section 2](#), is one reason to believe that the knowledge norm of assertion (or something like it), is correct.¹³

6 Deduction contexts

As discussed in [von Fintel & Gillies 2010](#), [Giannakidou & Mari 2016](#), [Lassiter 2016](#), and [Mandelkern 2016](#), *must* ϕ can appear in the conclusion of a deduc-

account leads to inconsistency. The first two components conspire to require ϕ to be known, which contradicts the last component, which requires ϕ not to be known.

¹³ Here is another argument for the knowledge norm, based on data from [Lassiter 2016](#):

- (i) [From [Lassiter 2016](#): p. 155-156]
 - a. I've seen these before and they must be aftermarket but I don't know for sure.
 - b. #I've seen these before and they are aftermarket but I don't know for sure.

Lassiter provides a few naturally occurring examples like (ia) demonstrating that an utterance of *must* ϕ can be accompanied by an explicit admission that ϕ is not known, and he claims that corresponding unmodalized assertions like (ib) are infelicitous. Under the picture I have sketched of the relationship between assertions of *must* ϕ and ϕ using the knowledge norm of assertion, the contrast between (ia) and (ib) is expected. If asserting ϕ is tantamount to claiming to know ϕ for sure, then following an assertion of ϕ with “but I don't know ϕ for sure” as in (ib) is predicted to be infelicitous, or as Lassiter puts it, “Moore's-paradoxical”. To assert *must* ϕ on the other hand is to claim to know for sure that ϕ holds throughout the *Best* worlds. Thus it is perfectly felicitous to follow up with the claim that one doesn't know ϕ for sure. That is, it is consistent for the speaker to claim both that she knows that ϕ holds in the *Best* worlds, but she does not know whether it holds in all the $\cap Epi$ worlds. If the knowledge norm of assertion were relaxed to, say, a truth norm (you can only assert ϕ if you believe it is highly likely that ϕ is true, [Matthew Weiner 2005](#)), then it is not clear why we get the contrast between (ia) and (ib), and in particular, why (ib) is infelicitous. After all, the speaker can believe that ϕ is highly likely to be true without knowing ϕ for sure. Thus the contrast between (ia) and (ib) is an argument in favor of the knowledge norm.

tion to ϕ . Assuming the *must* in such cases is epistemic, this is a challenge for the epistemic account. The reason is that if the premisses are known and they entail ϕ , then ϕ is also known. Thus the epistemic account predicts the *must* ϕ utterance to be infelicitous, contrary to fact. Consider a typical example of *must* ϕ in a deduction in (18):

- (18) [Adapted from Lassiter 2016]
 A teacher is explaining to a student that there is only one number that is both prime and even:
- a. If x is prime and even, then x is 2.
 x is prime ... x is even ...
 So, x must be 2.

The *must* ϕ statement in (18a) is felicitous, yet the known premisses entail ϕ . This runs counter to the predictions of the epistemic account.

Interestingly, one could take any of the examples from Section 2 in which the speaker could not felicitously say *must* ϕ , embed the *must* ϕ statement into an overt deduction context like that in (18a), and by doing so, the judgments switch from infelicitous to felicitous. For example:

- (19) Phil, the cook from (9), cannot felicitously say, “Dinner must be ready.” But suppose he is explaining to his daughter how it is that he knows for sure that dinner is ready. Phil says:
- a. If the chicken and peas are done and the table is set, then dinner is ready.
 The chicken and peas are done and the table is set.
 Therefore, dinner must be ready.

I have argued that the epistemic account predicts Phil’s *must* ϕ utterance to be infelicitous. This result is correct for (9), but not for (19). So the evidence is divided, some of it cuts in favor of the epistemic account, and some against it. Even a speaker like Billy who directly sees rain can felicitously say *must* ϕ in a deduction context. This result runs counter to the predictions of both the epistemic account and the indirectness account.

- (20) Billy from (3) cannot felicitously say upon seeing rain, “It must be raining.” But suppose she 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:

Must ϕ is felicitous only if ϕ is not known

- a. 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.

Both the epistemic account *and* the indirectness account predict Billy's *must* ϕ utterance to be infelicitous, which is correct for (3), but incorrect for (20). So the fact that *must* ϕ can appear in the conclusions of deductions is a challenge for both the epistemic account and the indirectness account.

These deduction contexts pose a general challenge to any account of the felicity conditions of epistemic *must*. Any account — epistemic, indirect or otherwise — needs to explain why *must* ϕ utterances that are infelicitous like (3) and (9) become felicitous in deduction contexts. The difficulty is that the modal bases do not seem to have changed from the non-deduction contexts to the deduction contexts, the agents seem to be reasoning with the same exact information. So it is not immediately clear how to account for the contrast.

Note an important intuitive difference between non-deductive uses of epistemic *must* and deductive ones. The former give rise to the weakness or non-confidence intuition that is responsible for Karttunen's problem. The latter do not. The *must* ϕ utterances in (18a) through (20a) are not intuitively weak or non-confident. Thus, any attempt to explain *must* ϕ in deductions within the confines of the epistemic account, or the indirectness account, or any account that seeks to explain the intuition giving rise to Karttunen's problem, will struggle to explain why the relevant intuition is missing in deductive uses of *must*. For example, if a theorist sympathetic to the indirectness account were to claim that *must* ϕ is felicitous in deduction contexts because they impose indirectness (in the form of overt premisses), then this produces a mystery: Why does indirectness give rise to weakness/non-confidence intuitions in canonical cases, but not in deduction contexts? If the source of the intuition is indirectness, then it should arise whenever indirectness is present.¹⁴

¹⁴ A theorist sympathetic to the epistemic account could try a similar move: *Must* ϕ is felicitous in deduction contexts because the speaker suspends their knowledge for the duration of the deduction, i.e., they pretend that they do not know ϕ . In response to this approach, we have to ask a similar question: If lack of knowledge is responsible for the weakness intuition, then why don't we intuit weakness in a deduction context where the speaker is behaving as if she does not know ϕ ? Perhaps such an approach can be made to work, but the contrasting intuitions would need to be explained.

All of the theories of epistemic modality that we have considered so far are designed to derive the weakness or non-confidence intuition. Thus, the fact that this intuition is missing from deduction contexts poses a real problem. Faced with this problem, I believe we need to seriously consider the possibility that the *must* ϕ utterances in conclusions of deductions are a flavor of modality distinct from epistemic modality.

Flavors of modality are distinguished by the modal bases they quantify over. I propose that while epistemic *must* quantifies universally over what is known as restricted by reasonable assumptions (the *Best* set), deductive *must* quantifies universally over the possibilities compatible with the overtly uttered premisses. Uses of deductive *must* have a couple of distinguishing linguistic features: There must be overtly uttered premisses, and those premisses are taken to be true. Since deductive *must* quantifies only over worlds compatible with true propositions, the prejacent will also be taken to be true. This explains why deductive *must* ϕ utterances do not give rise to a weakness intuition, unlike epistemic *must* ϕ utterances.

This analysis also explains why people who know ϕ like Phil and Billy can only felicitously use *must* ϕ in the deduction contexts in (19) and (20). They cannot use *must* ϕ in the contexts in (3) and (9) because neither epistemic *must*, nor deductive *must* are felicitous in those contexts. Epistemic *must* is infelicitous because ϕ is known. Deductive *must* is infelicitous because there are no overt premisses.¹⁵

Given that languages sometimes employ particular lexical items for particular flavors of modality (e.g., *allowed* is a deontic but not an epistemic possibility modal), further evidence for the existence of deductive *must* can be found in lexical items that are used in deduction contexts, but not canonical contexts of epistemic modality. For example, in Iberian Spanish the modal *debe de* is preferred for typical examples of epistemic *must* like those in

¹⁵ The idea of treating *must* in deduction contexts as a different flavor from epistemic modality is already present in Giannakidou & Mari 2016, who build on Giannakidou 1999. They argue that the *must* appearing in deductions is alethic (i.e., logical or root necessity). It quantifies over all possibilities and says that ϕ holds in all of them. However, the uses of *must* in the deduction contexts in (18) through (20) do not seem to quantify over *all* possible worlds, rather just the subset produced by conjoining the premisses, as I have argued. Moreover, it is not clear under this view why alethic *must* can be used by Phil and Billy in deduction contexts, but not in the contexts from Section 2. That is, why can't Phil or Billy just choose to use alethic *must* rather than epistemic *must* in the contexts in Section 2? These considerations led me to propose the analysis of deductive *must* in the main text.

Must ϕ is felicitous only if ϕ is not known

Section 2, while *tiene que* is preferred in deduction contexts. Here are two examples:

- (21) Hillary is in her office and sees falling rain out the window. She received an e-mail that morning saying that a hollywood movie would be filmed outside that day, and that if it didn't rain they would be making fake rain, though the fake rain isn't supposed to start until 5 pm. Hillary looks at the clock, which reads 4:50 pm.
- a. *debe de estar lloviendo*
must of be raining
'It must be raining.'
 - b. *#tiene que estar lloviendo*
has to be raining
'It must be raining.'
- (22) Billy from (3) cannot felicitously say upon seeing rain, "It must be raining." But suppose she 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:
- a. *#entonces debe de estar lloviendo*
then must of be raining
'Therefore, it must be raining.'
 - b. *entonces tiene que estar lloviendo*
then has to be raining
'Therefore, it must be raining.'

A reviewer for *Semantics & Pragmatics* has drawn my attention to similar facts in Portuguese discussed in Marques 2020:

- (23) The ball is in A, B, or C.
It is neither in A nor in B.
- a. *#Deve estar em C.*
must be in C
'It must be in C.'
 - b. *Tem de estar em C.*
has to be in C
'It must be in C.'

These examples are just an initial step toward the work needed on deductive *must*. The main point for present considerations is that 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. If the analysis of deductive *must* as a distinct flavor of modality can be maintained, then such examples would not form a part of the class of data to be explained by a theory of epistemic *must* to begin with.

7 Conclusion

In this paper, I have argued that epistemic *must* ϕ utterances are felicitous only if ϕ is not known. Combining this felicity requirement with Kratzer's (1991) account of epistemic *must* results in the prediction that reasonable assumptions play a necessary role in the inference to ϕ . By requiring the use of reasonable assumptions, the epistemic account inherits Kratzer's (1991) explanation of the weakness intuition that Karttunen (1972) identified in (1). Giannakidou & Mari (2016) implement this idea as a presupposition.

To distinguish the epistemic account from the indirectness account proposed by von Stechow & Gillies (2010), I introduced pairs of examples in Section 2 in which two speakers have identical perceptions, yet the felicity of their *must* ϕ utterances differ. I argued, not that the indirectness account cannot explain these asymmetries, but that it is not clear how the indirectness account could do more than stipulate an answer.

The epistemic account however is able to offer an explanation of the differences that distinguish the examples in Section 2. Intuitively, the agents who cannot say *must* ϕ know ϕ , while those who can do not. Theories of knowledge such as Lewis 1996 give us deeper insight into these intuitions.

Future work includes the issues discussed in Sections 5 and 6. More work is needed on the knowledge norm of assertion to determine whether the epistemic account can be fruitfully explained as a quantity implicature. The facts surrounding the proposed distinction between epistemic and deductive *must* require further exploration.

Finally, there is a remaining question about the connection between epistemic modality and evidentiality. von Stechow & Gillies (2010) argue for a deep connection between these two phenomena. On the other hand, I have argued that the indirectness account of epistemic modality leaves certain facts unexplained, and in its place I have argued that the felicity conditions of epis-

Must ϕ is felicitous only if ϕ is not known

temic modals depend on knowledge. On the face of it, it may seem that I have therefore argued against the connection between modality and evidentiality. However, this is not so as I have only discussed modals, and have not considered the empirical facts surrounding evidential morphemes, or any of the rich literature on the semantics and pragmatics of evidentiality (e.g., Willett 1988, Rooryck 2001a,b, Faller 2002, Matthewson, Davis & Rullmann 2007).¹⁶ Future work should include the systematic comparison of the behavior of epistemic modals and indirect evidentials in the kinds of contexts I have introduced. A reviewer for *Semantics & Pragmatics* suggests that, for example, Quechua evidentials may pattern with English epistemic modals in the contexts from Section 2. If future research were to bear this out for Quechua or any other language, this would suggest that there is indeed a connection between epistemic modality and evidentiality in at least some cases. However, if modals and evidentials pattern together in the contexts I have constructed, it does not mean that the indirectness account is correct, since it does not explain such patterns, as I have argued. Instead, the arguments I have made for modality would carry over to evidentiality: The use of evidentials would also seem to depend, at least in part, on knowledge.

References

- Blome-Tillmann, Michael. 2009. Knowledge and presuppositions. *Mind* 118 (470). 241–294. <https://doi.org/10.1093/mind/fzpo32>.
- Fagin, Ronald, Joseph Y Halpern, Yoram Moses & Moshe Vardi. 1995. *Reasoning about knowledge*. Cambridge, MA: MIT press.
- Fagin, Ronald & Joseph Y. Halpern. 1987. Belief, awareness, and limited reasoning. *Artificial Intelligence* 34(1). 39–76. [https://doi.org/10.1016/0004-3702\(87\)90003-8](https://doi.org/10.1016/0004-3702(87)90003-8).
- Faller, Martina T. 2002. *Semantics and pragmatics of evidentials in Cuzco Quechua*. Palo Alto, CA: Stanford University dissertation.
- Faller, Martina T. 2012. Evidential scalar implicatures. *Linguistics and Philosophy* 35(4). 285–312. <https://doi.org/10.1007/s10988-012-9119-8>.

¹⁶ Matthewson, Davis & Rullmann (2007) and Rullmann, Matthewson & Davis (2008) in particular argue for a connection between St’át’imcets evidentials and epistemic modality that is nevertheless very different in its details from the view in von Stechow & Gillies 2010. In particular, they argue that modals can either encode information about source or force, but not both.

- von Fintel, Kai & Anthony S. Gillies. 2010. *Must . . . stay . . . strong!* *Natural Language Semantics* 18(4). 351–383. <https://doi.org/10.1007/s11050-010-9058-2>.
- von Fintel, Kai & Irene Heim. 1997–2023. Intensional semantics. Ms. MIT. <https://github.com/fintelkai/fintel-heim-intensional-notes>.
- Frank, Michael C. & Noah D. Goodman. 2012. Predicting pragmatic reasoning in language games. *Science* 336(6084). 998–998. <https://doi.org/10.1126/science.1218633>.
- Franke, Michael. 2011. Quantity implicatures, exhaustive interpretation, and rational conversation. *Semantics and Pragmatics* 4(1). 1–82. <https://doi.org/10.3765/sp.4.1>.
- Franke, Michael & Tikitú de Jäger. 2011. Now that you mention it: Awareness dynamics in discourse and decisions. In Anton Benz, Christian Ebert, Gerhard Jäger & Robert van Rooij (eds.), *Language, games, and evolution: Trends in current research on language and game theory*, 60–91. Berlin, Heidelberg: Springer. https://doi.org/10.1007/978-3-642-18006-4_4.
- Geurts, Bart. 2010. *Quantity implicatures*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511975158>.
- Giannakidou, Anastasia. 1999. Affective dependencies. *Linguistics and Philosophy* 22(4). 367–421. <https://doi.org/10.1023/A:1005492130684>.
- Giannakidou, Anastasia & Alda Mari. 2016. Epistemic future and epistemic *must*: Nonveridicality, evidence, and partial knowledge. In Joanna Błaszczak, Anastasia Giannakidou, Dorota Klimek-Jankowska & Krzysztof Migdalski (eds.), *Mood, aspect and modality revisited: New answers to old questions*, 75–117. Chicago, IL: University of Chicago Press. <https://doi.org/10.7208/chicago/9780226363660.003.0003>.
- Goodhue, Daniel. 2016. Epistemic *must* is not evidential, it's epistemic. *North East Linguistics Society (NELS)* 46. 321–334. <https://semanticsarchive.net/Archive/jhlMzE5Y/>.
- Grice, Herbert Paul. 1989. *Studies in the way of words*. Cambridge, MA: Harvard University Press.
- Karttunen, Lauri. 1972. *Possible and must*. In John P. Kimball (ed.), *Syntax and semantics*, vol. 1, 1–20. New York, NY: Academic Press. https://doi.org/10.1163/9789004372986_002.
- Katzip, Roni. 2007. Structurally-defined alternatives. *Linguistics and Philosophy* 30(6). 669–690. <https://doi.org/10.1007/s10988-008-9029-y>.
- Kratzer, Angelika. 1991. Modality. In Arnim von Stechow & Dieter Wunderlich (eds.), *Semantics: An international handbook of contemporary research*,

Must ϕ is felicitous only if ϕ is not known

- 639–650. Berlin: de Gruyter. <https://doi.org/10.1515/9783110126969.7.639>.
- Lassiter, Daniel. 2016. *Must*, knowledge, and (in)directness. *Natural Language Semantics* 24(2). 117–163. <https://doi.org/10.1007/s11050-016-9121-8>.
- Lewis, David. 1996. Elusive knowledge. *Australasian Journal of Philosophy* 74(4). 549–567. <https://doi.org/10.1080/00048409612347521>.
- Mandelkern, Matthew. 2016. A solution to Karttunen’s problem. *Sinn und Bedeutung (SuB)* 21. 1–18. <https://doi.org/10.18148/sub/2018.v21i2.170>.
- Marques, Rui. 2020. Epistemic future and epistemic modal verbs in Portuguese. *Journal of Portuguese Linguistics* 19(1). <https://doi.org/10.5334/jpl.243>.
- Matthewson, Lisa. 2015. Evidential restrictions on epistemic modals. In Luis Alonso-Ovalle & Paula Menendez-Benito (eds.), *Epistemic indefinites: Exploring modality beyond the verbal domain*, 141–160. Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199665297.003.0007>.
- Matthewson, Lisa, Henry Davis & Hotze Rullmann. 2007. Evidentials as epistemic modals: Evidence from St’át’imcets. *Linguistic Variation Yearbook* (1). 201–254. <https://doi.org/10.1075/livy.7.07mat>.
- Ozturk, Ozge & Anna Papafragou. 2015. The acquisition of epistemic modality: From semantic meaning to pragmatic interpretation. *Language Learning and Development* 11(3). 191–214. <https://doi.org/10.1080/15475441.2014.905169>.
- Rooryck, Johan. 2001a. Evidentiality, part I. *Glott International* 5(4). 125–133. <https://doi.org/1887/14555>.
- Rooryck, Johan. 2001b. Evidentiality, part II. *Glott International* 5(5). 161–168. <https://doi.org/1887/14586>.
- Rullmann, Hotze, Lisa Matthewson & Henry Davis. 2008. Modals as distributive indefinites. *Natural Language Semantics* 16(4). 317–357. <https://doi.org/10.1007/s11050-008-9036-0>.
- Stalnaker, Robert. 1984. *Inquiry*. Cambridge, MA: MIT Press.
- Weiner, Matt. 2007. Norms of assertion. *Philosophy Compass* 2(2). 187–195. <https://doi.org/10.1111/j.1747-9991.2007.00065.x>.
- Weiner, Matthew. 2005. Must we know what we say? *The Philosophical Review* 114(2). 227–251. <https://doi.org/10.1215/00318108-114-2-227>.
- Willett, Thomas. 1988. A cross-linguistic survey of the grammaticization of evidentiality. *Studies in Language* 12(1). 51–97. <https://doi.org/10.1075/sl.12.1.04wil>.

Williamson, Timothy. 1996. Knowing and asserting. *The Philosophical Review* 105(4). 489-523. <https://doi.org/10.2307/2998423>.

Yalcin, Seth. 2018. Belief as question-sensitive. *Philosophy and Phenomenological Research* 97(1). <https://doi.org/10.1111/phpr.12330>.

Daniel Goodhue
daniel.goodhue@gmail.com