Must, knowledge, and (in)directness

Daniel Lassiter
Stanford University

This is a prefinal version of a paper in press at *Natural Language Semantics*. Please consult the *NLS* website for the authoritative version.

Abstract This paper presents corpus and experimental data that problematize the traditional analysis of *must* as a strong necessity modal, as recently revived and defended by von Fintel and Gillies (2010). I provide naturalistic examples showing that *must p* can be used alongside an explicit denial of knowledge of *p* or certainty in *p*, and that it can be conjoined with an expression indicating that *p* is not certain or that *not-p* is possible. I also report the results of an experiment involving lotteries, where most participants endorsed a sentence of the form *must not-p* despite being instructed that *p* is a possibility. Crucially, endorsement was much higher for *must* in this context than for matched sentences with knowledge or certainty expressions. These results indicate that the requirements for felicitous use of *must* are weaker than for *know* and *certain* rather than being at least as strong, as the epistemic necessity theory would predict. However, it is possible to account for these data while retaining the key insights of von Fintel and Gillies' analysis of the evidential component of *must*. I discuss several existing accounts that could be construed in this way and explain why none is completely satisfactory. I then propose a new model that embeds an existing scalar theory into a probabilistic model of informational dynamics structured around questions and answers.

Keywords: epistemic modality, evidentiality, knowledge, inference, probability

1 Introduction

It is widely assumed in philosophical circles that the English item *must* denotes an epistemic necessity operator. By contrast, in linguistic semantics it is—as von Fintel and Gillies (2010) note—a virtual "Mantra" that *must* is semantically weak in some sense. von Fintel and Gillies (2010) trace the Mantra back to Karttunen (1972), who writes:

> There is a striking difference between the logical necessity operator and words like *must*. Consider the two expressions in (1) and (2):
>
> (1) John must have left. (2) John has left.
>
> In any of the standard modal logics, $\Box p$ is a stronger expression than p. However, there is an inverse relation between [these] two sentences ... In stating (1), the speaker indicates that he has no first-hand evidence about John's departure, and

neither has it been reported to him by trustworthy sources. Instead, (1) 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. A man who has seen John leave or has read about it in the newspaper would not ordinarily assert (1), since he is in the position to make the stronger claim in (2). (Karttunen 1972: 11-12)[1]

As von Fintel and Gillies note, many other linguists have repeated this claim in various forms. The most influential version is that of Kratzer (1991), who formalizes the weakness intuition by making *must*, in effect, a universal quantifier over a maximally normal subset of the epistemically possible worlds ($\mathcal{E}$). Thus *must p* is compatible with there being $\neg p$-worlds in $\mathcal{E}$ as long as they are not maximally normal.

However, Palmer (1979: 59) points out that the intuitions supporting the "weakness" claim are confounded with indirectness of evidence. Indeed Karttunen hints at this—"... *no first-hand evidence* about John's departure ..." (my emphasis) —but Palmer's diagnosis is more explicit: "It is the notion of deduction or inference from known facts that is the essential feature of *must*, not just the strength of commitment by the speaker." Von Fintel and Gillies 2010 agree but go further, arguing that *must p* conveys indirectness of evidence but does not convey weakness of any kind: it is false if any $\neg p$-worlds are epistemically possible. They marshal an impressive variety of arguments for the "strength" claim, adding two novel formalizations of indirect evidential meaning which support the orthogonality of indirectness and weakness.

As an illustration of the empirical import of this debate, consider some examples of *must* from the message boards of the genealogy website Ancestry.com.

(3) I have learned that Flora & Dr Alexander had also a brother Dr Murdock MacLeod & a sister Margaret. Margaret married Kenneth MacLeod of Ebost on the Isle of Skye. Must have been about 1820.

(4) My great grandfather John HARTMAN married a lady by the name of Margaret KESSLER. John died in Colorado 1 Oct 1896. His wife was not in his will so must have died before that time. ... Does anybody know anything about this couple?

In (3), the example seems to convey that the conclusion marked by *must* was made on the basis of fragmentary historical records, and not, for example, an explicit report in an authoritative history book. In (4), the evidence is explicitly specified: the author infers that the wife must have died before the husband did, on the ground that she was not in his will. Ancestry.com users frequently provide an explicit specification of the evidence used to arrive at a *must*-conclusion (compare Stone 1994).

Is *must* "weak" or "strong" in these examples? This question is too coarse to be of much use, but we can break it down into a number of different strength-related questions. The most general involves the relationship between *must* and knowledge:

1 I've changed the example numbering, and updated the notation for the necessity operator from '*L*' to the now more common '$\Box$'.

- *Knowledge*: Do these examples commit their authors to the corresponding knowledge claim — for example, "His wife was not in his will, so I know that she died before that time"?

An uncontroversial entailment of knowledge is truth. If *must p* implies that *p* is known, then it should imply that *p* is true.

- *Veridicity*: Does (4) entail that the wife had in fact died — or only that this is a compelling inference given the available evidence?

Another plausible entailment of knowledge is certainty. Of course, this principle is controversial (indeed, quite implausible) if we are speaking of absolute certainty à la Unger (1971, 1975). But it is quite reasonable to suppose that knowledge implies the homely kind of certainty that is operative in ordinary, unmodified uses of *certain* and *sure* like the naturalistic examples in (5).[2]

(5) a. It's certain that Jimmy John's line-worker can't start a 2,000 store chain, but they could run off with some of the catering business.
b. This cannot happen unless both parties are sure that they will both mutually benefit from the arrangement in the long run.

A knowledge$\rightsquigarrow$certainty entailment would, for example, explain why example (6) sounds so strange (cf. Unger 1971).

(6) # He knows it's raining, but he isn't certain of it.

If this principle is correct, then an affirmative answer to the knowledge question above would also imply that *must p* should entail that *p* is certain.

(The knowledge$\rightsquigarrow$certainty entailment, while plausible, has been the subject of some controversy: see Sect. 7.2 below for discussion. In any case, many of the key arguments that I will give against the epistemic necessity theory go through even if the knowledge$\rightsquigarrow$certainty inference is not semantically valid, and I will try to be careful in what follows to make it clear which arguments require the validity of this principle as a premise.)

Furthermore, it is usually assumed that *certain* and epistemic *possible* are duals: *p* is certain if and only if *not-p* is not (epistemically) possible. This explains, for example, why (7) is plainly self-contradictory.

(7) It's not possible that it's raining, but it's not certain that it isn't.

2 Examples in this paper that are indicated as "naturalistic", "naturally-occurring", etc. were found using Google, except for Ancestry.com examples which were found using that website's search interface. I filtered out many false hits, non-epistemic examples, corrections, and examples which were plausibly produced by non-native speakers. (This, in addition to their casual style, largely accounts for the fact that many of the remaining examples come from online discussion forums, which often provide explicit discussion of origins or metadata indicating location.) I selected examples for discussion here based chiefly on (1) clear intuitive acceptability; (2) no indication that production error is involved (e.g., clear errors in the surrounding context, excepting minor spelling errors); (3) whether there was sufficient context in a small window around the example to render the intended interpretation clear; (4) whether the context rendered the relevant item unambiguously epistemic; and of course (5) relevance to the theoretical issue at hand. I have elected not to provide URLs since many of the URLs may not be stable. However, readers wishing to inspect the full context of an example can (as of this writing) usually recover it for most of the examples by copying them into the Google search window, with surrounding quotation marks.

The inferential connections between knowledge and certainty, and between certainty and possibility, would then correctly predict the oddity of both examples in (8).

(8) a. # He is certain that it's raining, but he considers it possible that it isn't.
b. # He knows it's raining, but he considers it possible that it isn't.

If these semantic connections hold, an affirmative answer to the knowledge question implies an affirmative answer to the certainty and impossibility questions:

- *Certainty*: Does (4) entail that its author, who explicitly acknowledges having very little evidence relevant to the conclusion drawn, is certain of the truth of *The wife died before that time*?
- *Impossibility*: Does (4) entail that its author considers the negation of the prejacent (*The wife had not died before that time*) impossible?

A "yes" answer to all of these questions would be in line with the classical treatment of *must* inherited from modal logic. In that account, *must* denotes in effect a universal quantifier over a set of worlds $\mathcal{E}$, containing the evaluation world, where all worlds in $\mathcal{E}$ are compatible with the relevant knowledge state (say, everything the speaker knows). Von Fintel and Gillies (2010) endorse this semantic theory, adding some further machinery to explain why *must* also indicates inference from indirect evidence (see Sect. 5.1 below). They endorse a "yes" answer to the veridicity question, and also to the impossibility question (see discussion in Sect. 4.2 below). While they do not explicitly discuss certainty, it is strongly implied at various points—and entailed by a positive answer to the impossibility question unless they reject the usual duality assumptions.[3]

This paper shows that the empirical facts are rather more complicated than this simple, attractive picture would suggest. Numerous corpus examples show that speakers can use *must* p when they do not know p, and are aware that they do not know p. Speakers use *must* p when they are not maximally certain of p; when they explicitly consider $\neg p$ to be a possibility; and when their stated grounds for concluding that p must be the case are less than fully compelling, and even explicitly stated to be so by the speaker. Examples (3) and (4) already begin to make this case, since the authors of these texts do not appear to be maximally confident of the truth of the conclusions that they are describing. Further empirical evidence for this claim is given in Sect. 2.

In addition, I present the results of an experiment which problematizes the claim that *must* implies knowledge and certainty. In a fixed context, a large majority of participants rejected both a knowledge claim and a certainty claim, while a clear majority accepted a matched *must*-claim.

3 Von Fintel and Gillies do note that *must* can sometimes be used to *convey* less-than-maximal confidence, but they argue that this effect is attributable to a secondary pragmatic inference involving the overall lower reliability of indirect evidence. They also note that *must*, along with uncontroversially strong expressions such as *There is no doubt at all*, can sometimes be used in a way that suggests lack of certainty via pragmatic reasoning about why someone would try so hard to persuade a listener to have maximal confidence (as in Hamlet's "The lady doth protest too much, methinks"). They rightly point out that this phenomenon should not be treated by weakening the semantics. However, consideration of the numerous examples in Sect. 2 suggests that the association between *must* and less-than-maximal confidence cannot be explained in this way in many cases. For instance, examples which explicitly indicate non-maximal confidence in the prejacent would not be well analyzed as overzealous efforts to persuade the listener to have maximal confidence in the prejacent.

Strikingly, most participants endorsed *must p* even while there was nearly universal agreement with *not-p is possible* in the same context. In addition to its implications regarding plausible entailment relations among these items, this experiment suggests that purely statistical evidence can license the use of *must*, even when this evidence is clearly insufficient for knowledge or certainty.

To be clear, this paper will not attempt to provide a conclusive answer to one of the core questions addressed by von Fintel and Gillies (2010)—whether *must p* entails *p*. My empirical claims center on the broader claim that *must* is well-analyzed as a strong epistemic necessity modal, convey knowledge of its prejacent. While I will argue that the arguments for veridicity are not compelling, nothing I will say rules out associating *must p* with some weaker evidential relation to *p* — as I will propose below — while also requiring that it cannot be true unless *p* is. (This is not to say that the implications of a weak but veridical *must* would be unproblematic: see Sect. 7.1 for a brief discussion with reference to theories of assertion.) In any case, the key claim of this paper is that, whether or not *must p* entails *p*, this entailment is not due to any simple relationship between *must* and knowledge.

While the empirical situation is complex, the evidence presented in this paper places important constraints on the correct theory of *must*. Most crucially, this theory must be able to make sense of the frequent use of this item to mark propositions believed on the basis of inductive inference, with high but non-maximal confidence. The final part of the paper surveys a number of existing accounts, pointing out their strengths and weaknesses, and proposes a new account built around a question-based probabilistic model of knowledge representation and discourse.

2 *Must* in the wild

This section discusses a number of corpus-derived examples illustrating the use of *must* and its interaction with *know* and other modals and confidence expressions. I will argue that the only way to make sense of these examples without attributing massive error to speakers is to allow that speakers can use *must p* when they do not know *p*, are not certain of the truth of *p*, and hold out the possibility of $\neg p$.

2.1 Lack of knowledge or entailing evidence

Genealogical discussions, such as those from which (3) and (4) above were collected, provide many informative examples of the use of *must*. These discussions are frequently devoted to speculation about the lives of unknown, long-dead persons, and they attempt to use fragmentary historical records, general reasoning, and hearsay to piece together details and dates of births, deaths, marriages, and migrations. Genealogists and other historians frequently use *must* to mark inferences which have been made in this fashion, despite clearly being aware that they do not know anything that entails the truth of their conclusions.

Indeed, it is well known among genealogists that the sources on which they rely are often unreliable (Mills 1999). If epistemic *must* were only appropriate when one knows or is certain of the conclusion drawn, we would expect that competent genealogists would almost never use *must*; and that we, knowing about the possible unreliability of sources, would perceive most instances of *must* in this context as infelicitous. But *must* is quite frequent in genealogical discussions, and seems to be a natural way to report reasonable, but fallible, inferences about past events.

Consider (9), taken from *The Plymouth Colony Archive Project*'s discussion of the fates of *Mayflower* passengers.

(9) Goodman, John. "Died soon after arrival in the general sickness," Bradford, p. 447. Goodman was still alive in mid-January 1621 (Mourt's Relation, pp. 45-48), although not in good physical shape. He is listed as one of those who received land in 1623 (PCR 12: 4). He is not listed among those who were part of the cattle division of 1627, so he must have died by then.

If *must* implies knowledge, this example should be felicitous only if the authors believe that they are in a position to know that Goodman died by 1627 — and so, inter alia, that there is no other explanation for Goodman's failure to appear on a certain list in 1627. That is, they would have to believe that they are in a position to know the following:

(10) The only way that a farmer who arrived on the *Mayflower* in 1620, is reported to have died "soon" afterwards, was sick in 1621, and was reported to be alive in 1623, could fail to be in a list of farmers in 1627 is that he died in the meantime.

But no modern author could reasonably self-ascribe knowledge of this proposition, direct or otherwise: the reasoning involves too many steps of inference and contains too many sources of uncertainty. Alternatively, we could suppose that the authors are uncooperative or flouting conversational norms for pragmatic effect. Neither explanation seems very plausible here.

A diagnosis that is more charitable to the authors of the *Plymouth Colony Archive Project* is that the use of *must* does not require knowledge: it merely requires very high confidence. That is, (9) does not convey that *Goodman had died by 1627* is the only possible explanation of the documentary evidence. Instead, the authors present Goodman's death as the *best explanation* of the evidence presented. It is indeed a good explanation, and it is rendered highly probable by what is reported; but it is in no way entailed by this evidence.

More examples like this could easily be adduced: for example, similar lessons emerge from the two genealogical examples that we saw already in Sect. 1. Example (4)—*His wife was not in his will so must have died before that time*—invokes a line of reasoning which is very plausible, but hardly foolproof. If *must* were an epistemic necessity operator, we would be forced to conclude, uncharitably, that the author of (4) does not realize that marriages can end for reasons other than death.

2.2 Explicit denial of (certain) knowledge

Perhaps the most direct argument against the analysis of *must* as an epistemic necessity operator is that speakers sometimes use *must* alongside an explicit denial that the speaker knows the prejacent with certainty. The following naturally-occurring examples illustrate.

(11) 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 <u>it must be 138,000</u>, but <u>I don't know for sure</u>.

(12) They don't repaint the helmets, I don't know for sure why, but it must be too much of a hassle to do.

(13) Spencer's "Native Soil" prints that I've seen are traditinal darkroom prints. I've heard he is currently working with inkjet printing which must mean PS work too, but I don't know that for sure.

(14) This is not non-stick. There must be others but I don't know for sure.

(15) In Despeado34's post the picture from JAXTRIUMPHGUY shows a nice little rack on the gas tank. I've seen these before and they must be aftermarket but I don't know for sure.

(16) I don't know for sure, sweetie, but she must have been very depressed. A person doesn't do something like that lightly.

(17) It's not a bad old thing. :) Must be '60s but I don't know for sure.

(18) Freeney must have called Josh, and Josh must have told me, and I must have been relieved. But I don't know for sure how I reacted, because I can't remember anything about the moment I learned that Ray was dead, and I didn't write a word about it in my journal. I must have thought that I could finally forget.

(19) Now this mortician had this great idea. He must have been Irish but I don't know for sure.

(20) [Where did the idea for the show "Bathtime Tuesday" spring from?] I think my dog was around and I looked over at the bathtub. It must have been a Tuesday (but I don't know for sure) and thought "Bathtime Tuesday!"

An interesting feature of these examples is that all of them involve the complex knowledge predicate *know for sure*, and sound less acceptable if *for sure* is omitted. For example, a modification of (13) with *for sure* omitted sounds fairly odd.

(21) ?# Inkjet printing must mean PS work too, but I don't know that.

In this light, it might appear that the epistemic necessity theory could deal with these examples by supposing that *must* expresses knowledge, but that knowledge is possible in the absence of certainty. (Recall that we are talking about ordinary uses of *certain*, not absolute certainty in the philosophical sense. Obviously knowledge is possible in the absence of the latter, assuming skepticism is false.) On this supposition, *know for sure* could be analyzed as a complex predicate expressing knowledge *and certainty*, in contrast to *know* which would express knowledge while leaving room for uncertainty. The examples just seen would then be compatible with the claim that *must* expresses (just) knowledge.

While this suggestion has some initial plausibility, it is not sufficient to rescue the epistemic necessity theory for at least two reasons. First, the examples are specifically designed to raise the possibility that the prejacent of *must* is false. The parenthesized material in (22), for example, clearly suggests that the speaker is unable to exclude the possibility that the day in question was not a Tuesday.

(22) It must have been a Tuesday (but I don't know for sure).

But, as von Fintel and Gillies (2010: Sect. 4.3) argue, the treatment of *must* as an epistemic necessity operator predicts that it should be pragmatically bizarre for someone to utter *It must have been a Tuesday* alongside a statement advising the listener that it is a serious possibility that it was not a Tuesday. I see no reason to doubt their theoretical analysis on this point. But unless it is flawed, (22) and the other examples in (11)-(20) give us strong reason to doubt that *must* is "strong" in the intended sense.

Second, the experimental evidence discussed in Sect. 3 below problematizes the strategy for retreat that we are considering. There I will show that, with the epistemic background held fixed, participants accept a prompt involving *must* much more readily than one involving *knows*. *Knows*, on the other hand, is accepted at the same rate as *certain*. These experimental results thus provide weak support for the assumption that knowledge entails certainty. More importantly, they reinforce the key implication of examples (11)-(20): *must p* can be true when *p* is not known. See Sect. 3 for further discussion.

While I think that the entailment from knowledge to (ordinary) certainty is reasonable—and I will continue to use it as a premise in some of my arguments—my arguments against the epistemic necessity theory of *must* do not stand or fall with this assumption. The pragmatic argument borrowed from von Fintel and Gillies (2010) above suffices to make the point: *must p* should be infelicitous alongisde any expression that functions to indicate that *not-p* is a live possibility. In addition, there are further empirical problems discussed below involving *possible*, an item whose semantical connection to *must* is not in doubt.

To the extent that knowledge *does* imply certainty, though, we need some other explanation of why the examples in (11)-(20) prefer reinforcement by *for sure*. One possibility is that the preference for *for sure* could be due to a pragmatic requirement that the strong contrast in these examples be explicitly marked, either by modification or by intonation. In support of this suggestion, note that even expressions of high but non-maximal confidence show a strong preference for modification in such contrastive contexts.

(23) Will Bill come to the party?
  a. It's almost certain that he will, but I don't know that for sure.
  b. ?? It's almost certain that he will, but I don't know that.

In addition, the examples involving *know* seem to improve markedly with contrastive focus on the verb.

(24) It's almost certain that he will, but I don't *KNOW* that.

(25) Inkjet printing must mean PS work too, but I don't *KNOW* that.

Regardless of whether the latter suggestion turns out to be correct, examples (11)-(20) are problematic for the epistemic necessity theory. If knowledge implies (ordinary) certainty, they directly contradict the epistemic necessity theory. If *must* implies knowledge, but knowledge does not require certainty, they still raise pragmatic issues that are quite awkward for the epistemic necessity theory. We have good reason to be skeptical of the claim that *must* expresses such an operator.

2.3 *Must* glossed with attenuated certainty expressions

In the following naturalistic examples, authors explicitly follow up a *must*-sentence with an expression suggesting uncertainty about the truth of the prejacent of *must*. To the extent that epistemic necessity implies certainty, these examples are again problematic for an analysis of *must* as expressing epistemic necessity.

(26) That face... was it not the same face as the shepherd he had met on the way? The one who gave him the straw? <u>It must be, he was almost certain</u>.

(27) That disturbing thorax, jointed into those long legs, casting a greater shadow than they have any right to. Yes, <u>it must be the legs. I'm almost certain</u>. It's in the name after all–Daddy Long Legs.

(28) This spot might be good fishing, I've always thought, though I haven't seen a soul out there trying. <u>The land must be private, I'm almost certain</u>.

(29) Sophia trailed off. Another epiphany exploded in her mind. "<u>The man must be Bren. I'm almost sure of it</u>. You can see the way he looks at her. . . like she's his whole world."

(30) I just deactivated the email routing, and same happens, the fax stays as "retrying routing". <u>This must be a bug, I'm almost sure!</u>

(31) <u>There must be some sort of limitation. I'm almost sure</u> in that I cannot redeem this free night voucher towards my $3k / night stay.

(32) <u>Res must be for "RESET", I'm almost sure</u> that's what that recessed button is for.

(33) You should have never come back. You hear footsteps on the stairs. <u>It must be him. It's almost certain that it is him</u>. You don't have the strength to turn your head. ...

(34) Those Stanley made in England knives <u>must be old stock as i'm almost sure</u> they aren't made here anymore.

(35) Help! I finally found the words for this song through the mudcat search. It has reference to Grit Laskin who <u>must be the writer as I'm nearly certain</u> Pete Seeger recorded it.

All of these examples should be markedly infelicitous if *must p* entails that *p* is certain. That is, while *It is/I am almost certain that p* does not strictly entail uncertainty in *p*, it is difficult to see why a speaker who is in fact certain of the truth of *p* would choose to use this sentence to describe their confidence in the prejacent. After all, there is a briefer, more informative gloss that is presumably available: *It is/I am certain that p*. We have no reason to think that the authors of the examples in (27)-(35) are being uncooperative in the Gricean sense. So, these examples provide strong evidence that speakers use *must p* in natural speech and writing despite not being certain of *p*.

Even if we were to suppose that *must p* does not imply *p is certain*, these examples would be problematic on the pragmatic grounds already discussed above. Von Fintel and Gillies 2010 use the supposed infelicity of *must p* alongside any acknowledgement of the possiblity of *not-p* as a key argument for a strong theory of *must*. However, the speaker's choice to use *nearly/almost certain* rather than simply *certain* in (26)-(35) has the pragmatic effect of raising the possibility of *not-p*—just as *I'm nearly finished with my homework* has the secondary effect of drawing attention to the small but presumably non-zero amount of unfinished homework.

2.4 Explicit denial of certainty; acknowledgement of alternative possibilities

The strong necessity theory would lead us to expect a similar, perhaps even stronger, infelicity in examples (36)-(39). Here, *must p* is conjoined with an expression which entails, rather than merely implying, that *not-p* is a possibility. Consider (36), for example, where an assertion of *must p* is immediately followed by a denial of certainty in *p*.

(36) I have an injected TB42 turbo and dont like the current setup. There is an extra injected located in the piping from the throttle body.. Must be an old DTS diesel setup but Im not certain. Why would they have added this extra injector?

Example (37), while less direct, expresses a similar sentiment:

(37) There's one missing pepper on the ground a few feet away. A closer look reveals it has been chewed by something. I wouldn't put it past that pesky blue jay to have teeth but then I think it is unlikely he does. It must be a squirrel. What else can get onto a second floor balcony?

The author of the blog post quoted in (37) has pretty shaky ornithological credentials, but this author's status as an expert speaker of English is not in doubt. *Must* is used to mark the proposition *It is a squirrel that chewed the pepper*, which is incompatible with a proposition previously indicated to be possible but unlikely (*The blue jay has teeth* and, by implication, *The blue jay chewed the pepper*). (38) and (39) make a similar point.

(38) I refuse to believe that this one game, Lost Planet 2 DX11, which was previously 100% stable remember, is crashing because my overclock is unstable .... It's not impossible, granted, but IMO it is highly unlikely. There *must* be some other cause.

(39) Given the way you are using the SSD the space you are 'losing' must either be temporary cache files or error logs. It is also possible that you might have iTunes configured so that your iOS backup or downloaded apps are going to the boot SSD rather than the hard drive so you'll want to see how you iTunes is configured.

These examples again indicate that *must p* is compatible with a lack of certainty in *p*. In addition, they clearly have the conversational purpose of raising the possibility that *p* is false, and so should be pragmatically bizarre if the epistemic necessity theory is correct.

2.5 Implicit statistical reasoning

Example (40)—pieced together from messages and replies on the Ancestry.com message boards—illustrates the relevance of uncertain assumptions and statistical inference to the use of *must*. The key point is Author 1's account in (40c) of his earlier choice to employ *must*: "I was only assuming", he explains, and then describes the reasoning behind his assumption.

(40) a. *Author 1*: [Y]our man Lazarus must have sustained injuries at [Buena Vista] by his death date. ...

b. *Author 2*: I check the killed and wounded list for the Battle of Buena Vista and Lazarus wasn't listed under killed and wounded.

c. *Author 1*: Curious. I was only assuming that since Lazarus is listed as dying on March 2, '47, that it was from wounds suffered the week prior (Buena Vista was fought Feb. 22-23). His unit (as with all volunteers at the battle) was certainly in the thick of it. Nevertheless, as we all know, disease took a heavier toll on the troops than actual enemy fire. In my research, though, when I see a death date that close to the battle date, I tend to think that wounds played a part.

Author 1's choice to use *must* in (40a) is not well accounted for by supposing that he thought he knew (directly or otherwise) that Lazarus had been wounded. After all, he explains in (40c), "I was only assuming ..." and "I tend to think ... ". In addition, Author 1 is an expert on this domain—he mentions elsewhere that "I have been researching the battle of Buena Vista for several years"—and he notes that "we all know" that disease was the primary cause of death. It seems unlikely that this author forgot that disease was a possible cause of death when formulating (40a).

The explanation continues with a description of implicit statistical reasoning:

- Most soldiers who died in the war died of disease [i.e., $P(\text{disease} \mid \text{died})$ is high].
- Also, most who died had not been wounded [i.e., $P(\text{wounded} \mid \text{died})$ is low].
- However, death *within a week of battle* is usually attributable to wounds [i.e., $P(\text{wounded} \mid \text{battle\&died})$ is high].
- Since Lazarus died within a week of battle, Lazarus must have been wounded.

Author 1 seems to have chosen *must* to indicate that battle wounds were highly probable—despite being well aware that there are other plausible explanations of the evidence. In other words, the use of *must* somehow indicates, as Stone (1994) argues, that Lazarus' having been wounded is the *best explanation* of the available evidence.

3 An experiment on lotteries

The experiment reported in this section directly tests the predictions of the epistemic necessity theory against the (non-)entailment relations that were motivated by corpus evidence in the last section. I use a lottery scenario, where a certain proposition *p* (that an individual with one ticket wins the lottery) has a known probability which is very small but clearly non-zero. Conversely, the proposition that the individual will not win has very high but non-maximal probability.

The two theories under consideration make very different predictions about patterns of response. The epistemic necessity theory holds that *must p* implies knowledge of and (we are assuming) certainty in the truth of *p*, as well as the impossibility of *not-p*. So, this theory predicts that *must p* should be endorsed at most as often as a knowledge or certainty claim involving *p*, and that it should never be acceptable alongside *not-p is possible*. In contrast, the "weak" theory that we began to motivate in the last section holds that *must* does not imply knowledge or certainty. So, it should be possible in principle to set up a scenario where *must p* is more acceptable than knowledge and certainty claims involving *p*, and that it should sometimes be acceptable alongside a statement indicating that *not-p* is possible.

In addition, the examples we saw in Sect. 2.4 are interesting for a subtler reason which is examined in this experiment. We saw *must p* used alongside a variety of expressions entailing or implicating the possibility of *not-p* — including *not impossible*, *uncertain*, *not certain*, and *possible*. Strikingly, *might* does not appear on this list. Could it be that *might* is in fact dual to *must* — so that *must p and might not-p* is a contradiction — but *possible* is weaker than *might*? If we suppose further that *certain* is the dual of *possible*, this would have the striking implication that the weaker item's dual — *certain* — must be stronger than the stronger item's dual — *must*.

A third question of interest is whether strong but inconclusive statistical evidence can license the use of *must*. The Lazarus example discussed in Sect. 2.5 suggests that it can, but a single example leaves open questions about the generality of the effect.

3.1 *Might* is stronger than *possible*

In Lassiter 2016: ch.6 I provide experimental evidence that *possible* is weaker than *might*. The experiment used a forced-choice task to investigate whether participants would agree with a modal statement about a lottery scenario — either "It is possible that x will win" or "x might win", where x was an individual who had one of 1,000 total tickets. Three conditions varied the distribution of tickets among other ticket holders: either all individuals had an equal number ("even"), a small number held most of the tickets ("skewed"), or a single individual held almost all of the tickets ("very skewed"). While this experiment was intended to investigate variation in responses depending on the distribution of tickets across individuals, an interesting secondary effect emerged: there was a strong main effect of whether *possible* or *might* was used in the prompt, which did not interact with the choice of condition. It appears that participants are more likely in such contexts to endorse "It is possible that x will win" than "x might win" regardless of their baseline propensity to endorse the *might*-statement.

Figure 1 presents the results of this earlier experiment. The main effect of the choice of modal (*might* vs. *possible*) was highly significant, $p < .005$ in a nested model comparison. I interpreted this pattern of results as evidence that *might* requires a higher level of confidence than *possible*. This would explain why, in every condition, participants were more reluctant to endorse the *might*-statement than the corresponding *possible*-statement.[4]

This interpretation has an important implication for the analysis of *must* and *certain*. On the

4 Rudin (2016) offers an interesting proposal to explain the fact that *might* and *possible* are frequently used to indicate *significant* probability, while allowing that their literal semantics requires only *non-zero* probability. The key idea is that, when $P(p)$ is only slightly greater than 0, *might p* and *p is possible* are literally true, but may be judged false if the difference is not greater than 0 to a degree that is relevant in context. This idea is interesting and plausible, but it encounters two kinds of empirical problems unless we also suppose that *might* and *possible* have different literal interpretations. First, *possible* and *might* behave differently in the lottery scenario, but there is no reason why the choice of lexical item would condition what counts as a "relevant" difference if their literal meaning is identical. Second, if their literal meaning is identical then the effects on these items' duals should be symmetric, so that—if *certain* and *must* both indicate maximal confidence—they should be affected by granularity in the same way. But, as we will see shortly, *must* is weaker than *certain*, as we should expect given that *might* is stronger than *possible*. Nevertheless, considerations involving imprecise use of scalar expressions such as those that Rudin discusses are surely relevant to the pragmatics of epistemic modals, and it would be reasonable to adopt Rudin's pragmatic account as a supplement to the semantic theory developed here.

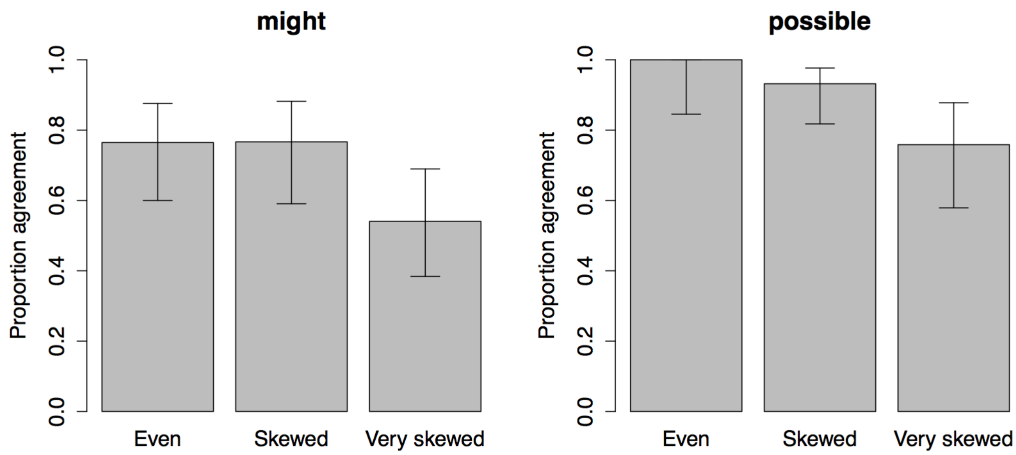

Figure 1 Results from the experiment reported in Lassiter 2016: ch.6 and summarized in the main text. The *y*-axis gives the proportion of participants who selected "Agree".

usual assumption that *must* and *might* are duals, it follows from the claim that *might* is stronger than *possible* that their respective duals have the inverse strength asymmetry. Unless we revise the usual assumptions about duality, we are forced to the conclusion that *might*'s dual *must* is weaker than *possible*'s dual *certain*.

These results, when combined with the evidence considered in Sect. 2, invite an explanation of the acceptability of examples like (38) and (39) along the following lines. While *must p and might not-p* would be a contradiction, *must p and it is possible that not-p* is consistent, and indicates high but non-maximal confidence in the truth of p. In contrast, if *must* indicated maximal confidence, it is difficult to see how we could render *must* and *might* duals without also making *might* and *possible* equivalent.

3.2 Materials and methods

The experiment reported here builds on the one reported in Lassiter 2016: ch.6, using a similar design, with an expanded range of items but without manipulating the ticket distribution.

After reading a background story, participants were asked to say whether they agreed or disagreed with a verbal prompt. The story described an individual, Bill, who held a single ticket in a raffle with 1,000 tickets in total. It was designed to highlight both the fact that there was a non-zero chance of Bill winning — indeed, he was as likely to win as any other individual — and also that his winning was highly unlikely. Here is the story that participants read (emphasis in original).

> Yesterday, Bill bought a single ticket in a raffle with 1000 total tickets. There were also **999** other people who bought **one** ticket each. That is, the tickets were distributed like this:

- People holding **one ticket**: Bill, Mary, Jane, ... [997 more]

The drawing was held last night, and the winner will be announced this evening.

The reason for putting the drawing in the past was to avoid possible interference from modal interpretations of the future, and also to ensure that all of the modals used could embed a negated proposition without awkwardness.

Participants saw a prompt chosen from the following list, and were asked to "Agree" or "Disagree" with the prompt after reading the scenario.

(41) a. Bill won the raffle. (*did*)
b. Bill did not win the raffle. (*did not*)
c. It is possible that Bill won the raffle. (*possible*)
d. Bill possibly won the raffle. (*possibly*)
e. We know that Bill did not win the raffle. (*know not*)
f. It is certain that Bill did not win the raffle. (*certain not*)
g. Bill certainly did not win the raffle. (*certainly not*)
h. Bill might have won the raffle. (*might*)
i. Bill must not have won the raffle. (*must not*)

570 Amazon Mechanical Turk workers were recruited and compensated for their participation. Each participant read the scenario alongside a single randomly chosen prompt, and chose to "Agree" or "Disagree" with the prompt. They were then asked to enter their native language, which was glossed as "The language that was spoken at home when you were a child". The experiment was run in four batches over the course of several months, as further conditions were added in response to theoretical developments. Some workers participated more than once, but the results reported include only their first responses: responses from repeat participants were filtered out in each subsequent batch.

3.3 Key predictions

I will assume a simple linking theory to bridge semantic theories with the behavioral data collected here: if p entails q, then q should be at least as acceptable as p relative to a fixed context which satisfies the presuppositions of both sentences.[5] We will therefore operationalize the theoretical construct of acceptability of a sentence in context as the proportion of English speakers who would choose "Agree" in that context, given the sentence as prompt. Note that this linking theory does not make any predictions about the relative acceptability of sentences that are not in a (contextual) entailment relationship.

5 This hypothesis makes sense on the assumption that, when p entails q, any competent user of the language who accepts p should also accept q unless pragmatic factors intervene. Previous experiments have shown that pragmatic factors can influence acceptability in a way that complicates these predictions: for example, an expression that gives rise to a false upper-bounding implicature might be accepted less readily than a stronger expression when both are true (e.g., Degen and Tanenhaus 2015, Nadathur and Lassiter 2014 among many others). However, I do not know of any reason to expect implicature or comparable pragmatic factors to interfere in the examples under consideration here.

This linking theory sets up a partial order based on entailment for each semantic theory of the items under consideration. The epistemic necessity account of von Fintel and Gillies (2010) and others predicts the relative acceptability ordering depicted in Figure 2. (Note that this ordering holds only when the evidential presupposition of *must* is satisfied — see Sect. 5.1. This is clearly the case in the vignette tested here. Note also that Figure 2 assumes that the epistemic necessity theorist wishes to endorse the knowledge⇝certainty inference. However, erasing this entailment does not affect the main conclusion: see the discussion in Sect. 3.4.)

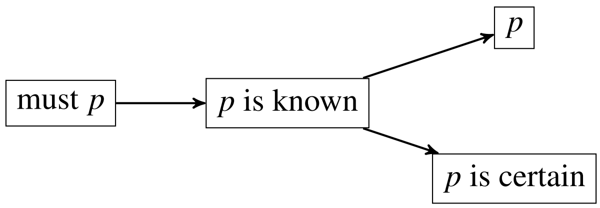

Figure 2 Partial ordering of some epistemic sentences in terms of entailment, as predicted by the analysis of *must* as an epistemic necessity operator. An arrow from *A* to *B* means "*A* entails *B*" or (given the linking hypothesis) "*A* is at most as acceptable as *B*, relative to a fixed context."

On the other hand, the corpus evidence presented in the previous section suggests a different picture: *must p* does not entail either *p is known* or *p is certain*. We also saw examples suggesting that it is roughly equivalent to *p is almost certain*, modulo the requirement of indirect evidence. If so, we predict that the partial ordering in terms of semantic strength should be either as in the left-hand or the right-hand panel of Figure 3, depending on whether we also assume veridicity.

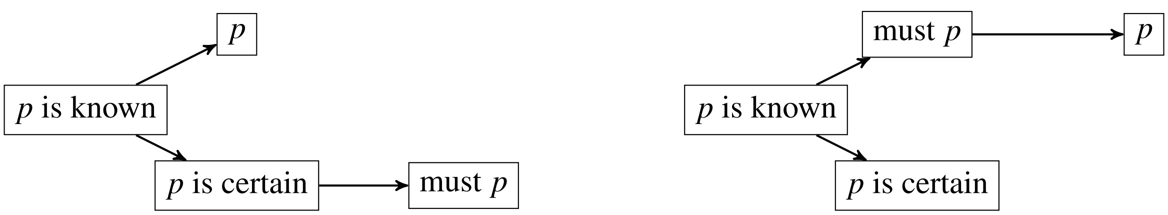

Figure 3 Two entailment orderings compatible with the corpus evidence and analysis presented in Sect. 2, differing only in that the ordering on the right treats *must p* as veridical (entailing *p*). Both orderings are valid only on the assumption that *must*'s indirectness requirement is satisfied.

A second set of predictions involve the comparison between *possible/certain* and *might/must*. If this experiment replicates the result from Lassiter 2016 discussed above, then *It is possible that Bill won* should be more acceptable than *Bill might have won*. If *might* asymmetrically entails *possible*, a slightly weaker prediction follows from our linking hypothesis: *p is possible* should be at least as acceptable as *might p*. As mentioned above, this entailment together with duality assumptions provides a second way to derive the prediction that *certain* is stronger than *must* (modulo the indirectness requirement). If so, our participants should be at least as willing to endorse *Bill must not have won* as they are to endorse *It is certain that Bill did not win*, and possibly more so.

As a side question included for interest, adverbial *certainly* and *possibly* also appeared among the prompts. These items have often been claimed to be more "subjective" than their adjectival counterparts, which have an "intersubjective" or "objective" flavor (e.g., Lyons 1977, Nuyts 2001). I conjectured that participants would be less strict in their use of subjective *possibly* and *certainly*, on the ground that a privately held certainty, or denial of possibility, may be felt to be less subject to public scrutiny and approbation if it turns out to be incorrect. If this is right, then participants should be more willing in the lottery scenario to endorse "Bill certainly did not win" than its adjectival counterpart "It is certain that Bill did not win." Similarly, they should be more willing to reject the subjective adverbial "Bill possibly won" on the basis of very low probability than the adjectival "It is possible that Bill won."

3.4 Results and discussion

I filtered out results from 14 participants who did not list English as a native language. I did not discard data from three participants who entered English along with one other language. This left 556 responses. Table 1 gives the numerical data, which is plotted in Figure 4 in ascending order by proportion of "Agree" responses.

Prompt	*did*	*know not*	*certain not*	*certainly not*	*must not*	*did not*	*possibly*	*might*	*possible*
Acceptance	0.07	0.21	0.25	0.54	0.58	0.69	0.74	0.80	0.92
Upper	0.13	0.32	0.35	0.67	0.70	0.81	0.84	0.89	0.98
Lower	0.02	0.11	0.15	0.42	0.46	0.56	0.62	0.71	0.85
N	71	54	57	79	57	65	62	53	58

Table 1 Numerical results by prompt. Upper and lower bounds refer to 95% bootstrap confidence intervals. See (41) for prompt label codes.

Assuming the linking hypothesis, these results are incompatible with the analysis of *must* as an epistemic necessity operator entailing knowledge of its prejacent (Figure 2). However, they are consistent with the analysis of entailment patterns given in section 2, summarized in Figure 3.

Following the entailment relationships depicted in Figure 2, the strong necessity theory of von Fintel and Gillies (2010) predicts the following acceptability relationships among experimental prompts. (I write "$A \succcurlyeq B$" as shorthand for "Prompt A is at least as acceptable as prompt B in the experimental scenario" — derived from entailment relations using the linking theory above. Note that the variable "p" in Figure 2 is being instantiated here as the negative sentence *Bill did not win the raffle*.)

(42) Relative acceptability predictions of strong necessity theory:

a. *did not* $\succcurlyeq$ *must not*

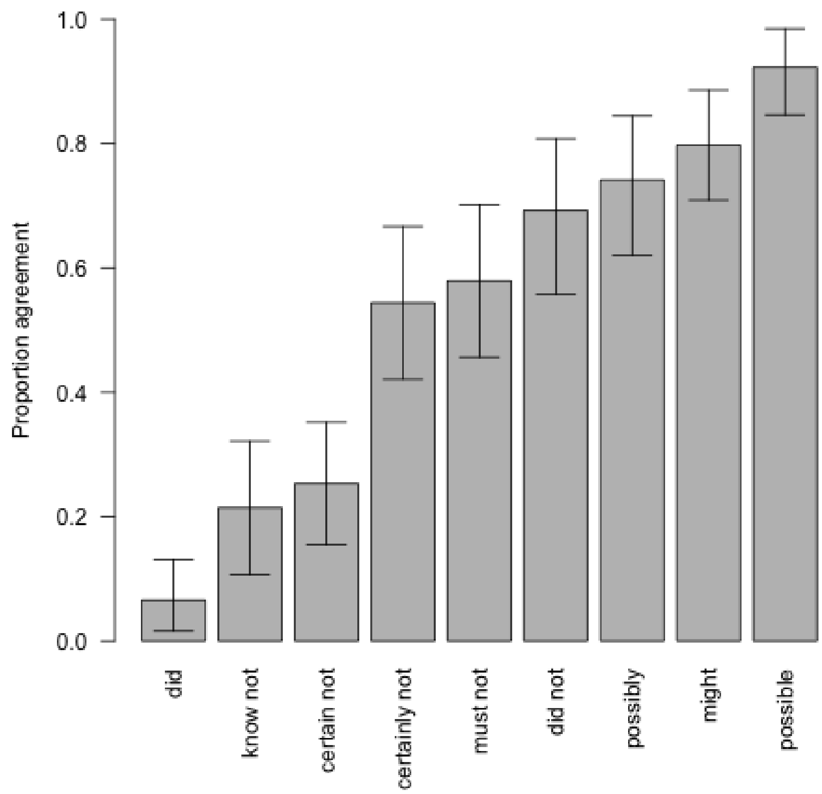

Figure 4

Results from the experiment described in the main text. The labels on the x-axis correspond to the prompts in (41). Error bars indicate 95% bootstrap confidence intervals.

b. *did not* ≽ *know not*

c. *certain not* ≽ *must not*

d. *certain not* ≽ *know not*

e. *know not* ≽ *must not*

Of these predictions, the first, second, and fourth are consistent with the experimental results. However, the predictions in (42c) and (42e) are strongly disconfirmed: contrary to prediction, the *must not* prompt was accepted much more often than either *certain not* or *know not*. Specifically, participants were far more likely to accept *Bill must not have won* than *We know that Bill did not win* (.58 vs. .21, $p < .0001$).[6] Likewise, participants accepted *Bill must not have won* more readily than *It is certain that Bill did not win* (.58 vs. .25, $p < .0005$).

These results would be difficult to explain if *must p* entailed *p is known*, *not-p is not possible*,

6 All p-values reported in this paper use a two-tailed bootstrap resampling test based on 50,000 samples.

and/or *p is certain*. Indeed, more than half of our participants were willing to endorse *Bill must not have won* in this scenario, even while they were generally unwilling to judge *Bill did not win* to be certain or known, and they assented almost universally to *not-p is possible*. These result militate strongly against the analysis of *must* as an epistemic necessity operator.

The alternative set of entailment relations suggested by the discussion in Sect. 2, and depicted in Figure 3, would make the following predictions.

(43) Relative acceptability predictions common to both "weak *must*" accounts in Figure 3:

a. *did not* ≽ *know not*

b. *certain not* ≽ *know not*

c. *must not* ≽ *know not*

All of these predictions are consistent with the experimental results. In particular, the entailment orders in Figure 3 predict that *must not* should be at least as acceptable as *know not*, and at least as acceptable as *certain not*. These are, recall, the experimental results noted above as problems for the epistemic necessity account.

If *must* is non-veridical (Figure 3, left), then we predict the additional relative acceptability relationship:

(44) *must not* ≽ *certain not*

This prediction is compatible with the experimental results, and indeed (as we already saw) *must not* was accepted at a much higher rate than *certain not*.

If we suppose instead that *must* is veridical (Figure 3, right), we make the prediction in (45) instead.

(45) *did not* ≽ *must not*

This prediction, too, is compatible with the experimental results: *must not* and *did not* appear to be roughly equally acceptable. While *did not* was accepted at a numerically higher rate, this difference was not reliably greater than zero ($p \approx .61$).

These results do not help us to decide whether *must* is veridical, then. Each version of the weak-*must* theory makes one correct prediction about which the other is silent. Neither makes a positive prediction which is disconfirmed. The results are thus compatible in principle with both veridicity and non-veridicity.

Several further features of the experimental results are worth noting briefly. First, the results are consistent with the supposition that knowledge implies certainty on the part of the knowing subject. While a finding of no difference does not of course prove that this entailment exists, the results do block one possible escape route for a defender of the epistemic necessity account, which we discussed briefly already in section 2.2 above. Specifically, given that *Bill must not have won* was found to be more acceptable than *It is certain that Bill did not win*, one might consider the possibility that *must* expresses knowledge, but that knowledge does not require certainty. While this much is compatible with the experimental results, the fact that *must not* was also more acceptable than *know not* indicates that this move would do little to rescue the epistemic necessity theory.

Second, the results confirm the prediction that *possible* should be at least as acceptable as *might*, which (we hypothesized) asymmetrically entails it. Indeed *possible* was strictly more acceptable, as expected (.92 vs. .80, $p < .05$). This replicates the previous finding that *possible* is weaker than *might* (Lassiter 2016: ch.6). The reliability of the effect in the present experiment was lower than in the earlier experiment, and in fact correcting for multiple comparisons using, say, a Bonferroni correction would render the effect marginal. However, this difference is likely attributable to two factors. First, the current experiment employed only the "Even" condition from the previous experiment, which had the highest baseline acceptance rate for both items. This could easily lead to a ceiling effect, where small but real differences in true strength are more difficult to detect at the extremes. In contrast, the earlier experiment found a highly significant main effect across several conditions, including two with lower baseline acceptance and thus a smaller ceiling effect. Second, the present experiment had less statistical power: the earlier experiment looked for main effects of the choice between *possible* and *might* across nearly twice as many data points (195, vs. 111 total in these two conditions in the current experiment).

A third, related prediction was derived from reasoning about duality relationships: if *possible* is weaker than *might* when all presuppositions are satisfied, then *certain* should be stronger than *must* in the same context. This predicts that *certain* should be at most as acceptable as *must*, and — as we have already seen — the results confirm this prediction (.25 vs. .58, $p < .0005$).

The final set of predictions is mostly orthogonal to the others, but is described here for completeness. I included *certainly* and *possibly* among the items out of interest, based on a hunch that participants would be more willing to use the more "subjective" adverbs in a risky way than their "intersubjective" or "objective" adjectival counterparts. The experiment confirmed this prediction. Participants were much more willing to reject *Bill possibly won* than *It is possible that Bill won* (.26 vs. .08, $p < .01$), and they were much more willing to accept *Bill certainly did not win* than *It is certain that Bill did not win* (.54 vs. .25, $p < .001$). The confidence associated with the use of a subjective certainty expression was lower, and the confidence associated with the use of a subjective possibility expression was higher, consistent with the suggestion that participants should be more willing to take risks in offering these judgments when the modal had a subjective flavor.

Summarizing the discussion so far, the results of this experiment are highly problematic for a theory according to which *must p* entails that *p* is known to relevant subjects. The *must not* prompt was found to be much more acceptable than the *know not* prompt in the lottery scenario that was tested. As expected, the results also problematized the claim that *must* is strictly stronger than *certain*.

The use of a lottery scenario in this experiment suggests some further questions about the relationship between *must*, confidence, and (merely) statistical evidence. Our participants accepted the use of *must p* in a situation in which *p* was known to have probability .999 — a very high probability, but one which clearly leaves room for *not-p*. Indeed, our participants almost universally endorsed a prompt of the form *not-p is possible*, even while more than half who saw a prompt of the form *must p* agreed. This pattern of responses has several important implications.

First, this finding corroborates the earlier conclusion, based on corpus evidence, that *must p* can be used alongside *not-p is possible*, *p is uncertain*, and a variety of other expressions indicating mild uncertainty about the truth of *p*. This is important because, as we will see below, von Fintel and Gillies' (2010) arguments for the epistemic necessity theory have emphasized that *must p* should never be acceptable alongside *any* indication of the possibility of *not-p*, no matter how weak.

Second, the fact that a majority of our participants agreed with *Bill must not have won* when there was merely statistical evidence for this conclusion shows that statistical considerations can be relevant to the use of *must* as long as they engender sufficient confidence. This supports the analysis of the Lazarus example in Sect. 2.5 above, where we suggested that *must p* indicates (merely) that the available evidence strongly favors the conclusion that p. In contrast, the relevance of strong but inconclusive statistical evidence to the use of *must* would be difficult to explain if *must* denoted an epistemic necessity operator. This consideration will play a major role in the eventual development of an alternative proposal in Sect. 6 below.

Third, a majority of the participants endorsed the description "Bill must not have won the raffle", even while Bill's prospects were indistinguishable from those of the other 999 ticket-holders. This is striking because the same participants would presumably, in the same situation, have been equally willing to endorse "Mary must not have won the raffle", "Jane must not have won the raffle", and so on for each of the remaining ticket-holders. Still, they probably would not have been willing to endorse "It must be that no one won the raffle", even though the prejacent of this *must*-claim is equivalent to the conjunction of "Bill did not win", "Mary did not win", etc. This reasoning suggests that *must* may have a meaning that is not closed under conjunction: *must p* and *must q* do not always imply *must* $(p \wedge q)$. See Sect. 6.2 below for further discussion and some important qualifications.

First, however, we should pause to reflect on some further arguments that have been offered against a "weak" theory along the lines of Figure 3 and in favor of a "strong" theory which encodes the entailments in Figure 2. As we will see, these arguments are of varying quality, and the most convincing of them are compatible with the "weak" theory that I am developing.

4 Countering some arguments for a strong theory of *must*

Von Fintel and Gillies (2010; henceforth "vFG") marshall a number of intuitive arguments in favor of a strong semantics for *must*, where the latter item entails that is (a) p is true, (b) p is known, (c) *not-p* is not possible, and (d) p is certain. Since our corpus and experimental investigations have led to a very different conclusion about points (b), (c), and (d), it is worth revisiting vFG's arguments now. I will argue that, of their four main arguments,

- the first is strictly irrelevant;
- the second is directly contradicted by the empirical data presented in Sects. 2-3; and
- the third and fourth provide some evidence for point (a), veridicity, but this is not sufficient to motivate the full package of (a-d), including knowledge, impossibility, and certainty entailments.

I will also argue that the third and fourth arguments are inconclusive even as regards veridicity. The success of these counter-arguments is not critical for the main point of the present work: the epistemic necessity theory is falsified by the failure of the knowledge, impossibility, and certainty entailments. However, I will take the time to problematize vFG's arguments for veridicity because—for reasons discussed in Sect. 7.1 below—a veridical version of the weak *must* interacts with theories of assertion in a way that some would consider undesirable.

4.1 "*Must* is not always weak"

vFG emphasize the distinction between indirectness and weakness: conclusions derived from indirect evidence can be maximally strong, and *must* can be used to signal indirectness without even hinting at uncertainty—for instance, in proofs.

(46) x is prime. x is even. Therefore, x must equal 2.

vFG attempt to use this observation to argue that *must* is not semantically weak. To my knowledge, though, no "Mantrista" has claimed that *must entails* a lack of certainty. For example, Kratzer (1991) treats *must p* as true whenever p is true in all of the maximally normal worlds among the epistemically possible worlds $\mathcal{E}$. However, Kratzer's semantics is silent on p's status in non-maximal worlds.[7]

vFG note that Portner (2009) responds in this way, and they reply that the Mantra is presumably not intended to apply to only *some* uses of *must*. But this is beside the point. The question is whether, if *must p* is always semantically weak, its use should always be associated with an inference that the speaker is uncertain about the truth of p. vFG assume that this should follow, but it is simply not so—any more than the fact that *some* is always semantically non-maximal leads inexorably to the prediction that it should always lead to an inference that *all* would be false. Indeed, *some* is frequently used in contexts in which the *not all* inference arises weakly or not at all (Degen 2015).

We can, however, extract a related argument against a Kratzer-style weak theory. Suppose that *must p* entails only that p is true in the maximally normal worlds in $\mathcal{E}$. Then, if English has some modal operator which requires that its propositional argument be true in all worlds in $\mathcal{E}$, *must p* should be associated with a quantity implicature to the effect that p is not true in all of the worlds in $\mathcal{E}$. Suppose that *certain* is such an expression. Then *It must be raining* should sometimes (though by no means always) implicate *It's not certain that it's raining*. If this is a good theory of *certain*, the fact that (46) does not tend to give rise to an uncertainty implicature may be a problem for Kratzer's semantics.[8]

However, the argument from implicature does not apply to a theory which combines weakness with an evidential component of the type that vFG themselves argue for. (We will discuss several possible implementations of such a theory in section 5-6.) The problem would return if there were a word *certain'* expressing both indirectness *and* universal quantification over $\mathcal{E}$; but the only plausible candidate is *must* itself, as analyzed by vFG. Since the existence of such an item is exactly what is in dispute, the occasional lack of a pragmatic inference to lowered confidence tells us nothing. On the other hand, if English allows us to express maximal confidence (using *certain*) or high-but-not-necessarily-maximal confidence with an evidential signal (using *must*), a speaker's choice to employ *must* could be explained by a desire to emphasize indirectness of evidence, even under full confidence.

7 With vFG, I'm assuming that *must* in (46) is an epistemic modal. See however Giannakidou 1999 for an argument that this use represents a separate, "alethic" sense of *must*.

8 Spoiler: this is not a good theory of *certain*. See Lassiter 2016: Sect. 5 for discussion of several better theories.

4.2 *Must* and weak epistemic operators

vFG observe that (47) is intuitively contradictory, and use this observation to argue that *must* entails maximal confidence.

(47) # It must be raining, but perhaps it is not raining.

If *perhaps* expresses existential quantification over $\mathcal{E}$, the oddness of (47) is unexpected according to weak-*must* theories, since these hold that *must p* does not exclude the presence of $\neg p$-worlds in $\mathcal{E}$. However, if *must* expresses universal quantification over $\mathcal{E}$, the infelicity of (47) is expected.

As stated, this argument misses the target: a weak theorist who wishes to maintain that (47) is a contradiction could simply adopt the usual assumption that *must* and *perhaps* are duals. If they are, (47) is ruled out as infelicitous no matter what *must* means, since it expresses the conjunction of *must p* and *not must p*.

vFG acknowledge this point briefly in a footnote, responding as follows: "There *are* strong necessity epistemic modals. So pick one and take its dual (e.g., *there is a vanishingly small chance that*). It'll be horrible when paired with *must* in examples like [(47)], we promise" (p.365, fn. 25). But in fact we saw several naturally-occurring examples with a similar form in section 2. Two of these examples are (48) and (49), repeated from (36) and (38) respectively. These examples closely resemble the frame that vFG claim is impossible.

(48) I have an injected TB42 turbo and dont like the current setup. There is an extra injected located in the piping from the throttle body.. <u>Must be an old DTS diesel setup</u> but <u>Im not certain</u>. Why would they have added this extra injector?

(49) I refuse to believe that this one game, Lost Planet 2 DX11, which was previously 100% stable remember, is crashing because my overclock is unstable .... <u>It's not impossible</u>, granted, but IMO it is <u>highly unlikely</u>. There <u>*must*</u> be some other cause.

On fairly uncontroversial assumptions, (48) and (49) are counter-examples to the claimed relationship between *must* and weak epistemic operators. The only way out would be to claim (for (48)) that *I am uncertain that p* does not commit a speaker to the possibility of *not-p*, and (for (49)) that *p is possible* is not entailed by *p is not impossible*. Neither response looks very promising.

In addition, the experimental results reported in Sect. 3 above indicate that participants are sometimes simultaneously inclined to endorse *must p* and *not-p is possible*. Specifically, 58% of participants agreed with *Bill must not have won* in the raffle scenario, while 92% agreed with *It is possible that Bill won*. Since participants had been randomly assigned to conditions, this result suggests that half or more of our participants were simultaneously inclined to endorse a claim of the form *must not-p* and a claim of the form *p is possible*.

None of this tells us about the status of *perhaps*. It could be that vFG are correct that (47) is contradictory. If so, *perhaps* is stronger than *possible*. Alternatively, it could be that *perhaps* is equivalent in strength to *possible*, and that vFG's intuition about (47) is not reliable. Either way, the evidence fails to show that *must* is maximally strong.

Finally, it is worth noting that the clash that we are observing in this case between intuitions and data derived from corpora and experiments suggests that we should not rely too heavily on intuitions in this domain. Multiple lines of evidence suggest that an initially plausible intuition —

that *must p* is incompatible with any indication that *not p* is possible — was not reliable. The same holds of the lottery scenarios, where empirical evidence contradicted the intuition that *Bill must not have won* should be categorically unacceptable.

4.3 Inferencing with conditionals

Speaking of intuitions, consider the following argument:

(50) a. If p, must q.
b. p.
c. $\therefore q$.

This has the feel of a valid argument. According to vFG, this intuition provides evidence for one plank of their account, namely, the veridicity of *must*. The reason is that, if *must p* did not entail p, the argument would be logically invalid. For example, on Kratzer's (1991) semantics, *If p, must q* is true as long as q is true in all of the maximally normal worlds in $\mathcal{E} \cap p$. But the actual world might be non-maximal, and so the fact that p is true in the actual world is consistent with $\neg q$ being true there as well.

If *must* is veridical, we have a clear explanation for the apparent validity of (50). Indeed, as I emphasized above, veridicity of *must* is formally compatible with the account being developed here (compare Figure 3). Since our rejection of the strong necessity theory of *must* did not require taking a stand on veridicity, we could stop here and not bother arguing against it. However, this would be an unsatisfying conclusion for those who share the intuition of Karttunen (1972) and Kratzer (1991) that *must p* does not entail p. In addition, as we will see in Sect. 7.1 below, a veridical but weak theory of *must* interacts in a possibly awkward way with theories of assertion; but no complications arise if we suppose that *must* is non-veridical. For these reasons, I will take a moment to sketch out a way to account for the feeling of validity that attaches to (50) within a theory that treats *must* as encoding only high confidence and indirectness. Note, however, that nothing in the rest of the paper really hinges on the success of this alternative suggestion (and indeed the rest of this section can be skipped without loss of continuity).

One way to avoid embracing veridicity in light of (50) is to note that the empirical basis of the argument involves intuitions about argument strength, and such intuitions are the subject matter of a well-developed body of psychological theorizing that may help in their interpretation. While linguists, philosophers, and psychologists alike have traditionally assumed that intuitions about argument strength come in some small number of discrete flavors — valid/invalid, say, or valid/invalid-but-consistent/contradictory — there is an increasing consensus in the psychological literature on reasoning that argument strength intuitions form a continuous scale. Typically, this scale is thought to be closely related or identical to the conditional probability of the conclusion given the premises (Oaksford and Chater 2007; Over 2009). (A closely related idea has been influential in the philosophical literature on conditionals and vagueness under the name "probabilistic validity": see Adams 1975; Edgington 1995, 1997.)

A key implication of this psychological theory of argument strength is that invalid arguments may sometimes be felt as extremely strong — or, if you like, as "intuitively valid" — as long as the truth of the premises guarantees that the conclusion will have high probability. Along similar lines, we might suppose that the sense of validity that seems to accompany (50) is really a sense that, no matter what the background context is like, the conclusion is very strong if the premises are true. (Indeed, if we identify argument strength with conditional probability this much would follow from the theory developed in Sect. 6 below, together with some plausible assumptions the meaning and use of conditionals.) Notably, this account would predict that logically valid arguments should always be maximally strong, since their conclusions have conditional probability 1 given the premises whenever it is possible to condition on the premises at all. Likewise, arguments whose conclusions contradict their premises are guaranteed to be maximally weak. This account thus does justice to the intuitions of strength that accompany logically valid arguments, without supposing that these intuitions of strength are indexing logical validity exclusively.

There is some disagreement in the psychological literature about whether people have access to a separate intuitive sense of argument strength which tracks strictly logical validity, in addition to the scalar notion of argument strength just described. Some arguments to this effect have been made, based in part on the effects of manipulating epistemic vocabulary in the conclusion of arguments (Rips 2001, 2002; Heit and Rotello 2010). However, follow-up work suggests that these effects can be explained in terms of the semantics of the epistemic items involved, and adduced additional empirical arguments that favor a unified scalar theory of argument strength intuitions (Lassiter and Goodman 2015b). We have every reason to suspect that the argument strength intuitions that kick in when we consider arguments with the form in (50) are tracking something correlated with, but qualitatively different from, logical validity.

To illustrate, consider a concrete instantiation of the argument form in (50). The first premise is based on (35), one of the corpus examples used above to problematize the claim that *must* entails maximal confidence.

(51) a. If Grit Laskin is on the record, he must be the writer (as I'm nearly certain Pete Seeger recorded it).
b. Grit Laskin is on the record.

Based on the psychological work cited above, we can reasonably expect that most participants who have been confronted with the premises in (51) would endorse the conclusion in (52).

(52) So, Grit Laskin is the writer.

However, as we discussed above (Sect. 2.3) the continuation in (51a) strongly implies that there is a slight possibility that Pete Seeger was *not* the recording artist. That is, the author is suggesting that the explanation for Laskin's appearance on the record *could* be that he was actually the recording artist, though this is very unlikely. The truth of the premises in (51) thus fails to render the conclusion indubitable, but the argument is very strong nevertheless — perhaps even "valid" in the intuitive sense. While this suggestion awaits empirical verification, I predict that an experimental investigation could uncover many such instances of arguments with the form of (50), where the conclusion is judged to be very strong even when we explicitly introduce mild uncertainty about whether q will be true if p is. If this is the case, we have an explanation of why these arguments "feel" valid which does not rely on the assumption that *must* is veridical.

None of this shows that the argument in (50) is *not* valid, of course. All it shows is that we *could* account for the feeling of validity that accompanies (50) in a well-motivated way without assuming that *must* is veridical.

4.4 Distancing

A fourth argument that vFG give in favor of their theory of *must* involves whether it is felicitous for someone who asserts *p*, *might p*, *must p*, etc. to avoid blame if *p* turns out to be untrue. vFG argue that such distancing is not acceptable with *must*, illustrating with the dialogue in (53).

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

Alex's response could, however, be appropriate with *might*, *ought* or *likely*.

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

In vFG's judgment, "saying *I only said it must be raining* is as bizarre as *I only said I ate all of the cookies*" (von Fintel and Gillies 2010: 366).

The argument seems to have two different points. First, since *I only said must p* is awkward, *must* indicates maximal confidence; else, *only* and *must* would "combine like old friends" (*ibid*). Second, vFG are concerned to show that *must* is veridical: otherwise, they argue, we could not explain the intuition that Alex said something false in (53), and that this is why she cannot back off and claim not to have taken a stand on the matter.

Beginning with the strength claim: we have already seen considerable corpus and experimental evidence indicating that *must* does not entail maximal confidence, and (53) does little to endanger this conclusion. If the logic of the "old friends" argument were valid, the dialogue in (53) should be fully acceptable when *must/might* is replaced by any item that is uncontroversially non-maximal, like *virtually certain*, *practically indubitable*, or *99.9% certain*. Contrary to this prediction, (55) sounds strange too:

(55) a. *Alex*: It's 99.9% certain that it's raining.
b. *Billy*: [opens curtains] No it isn't. You were wrong.
c. *Alex*: ?? I was not! Look, I didn't say it *was* raining. I only said it was *99.9% certain* that it was. Stop picking on me!

To my ear, the same oddity results if the epistemic expression is *virtually certain* or another expression of very high confidence. The use of these expressions seem to generate intuitions of blameworthiness, even though everyone agrees that they indicate non-maximal confidence and are not in semantic contradiction with the facts. If vFG are right that (53) is categorically unacceptable, we can only conclude that the ability to combine felicitously with *only* in these examples is not a reliable test of whether an expression implies maximal confidence.

Turning to the veridicity claim, the situation is again complex — and again, it might be a good idea to duck out and simply concede that *must* is veridical. Unwisely, perhaps, I will attempt a response here too. While the oddity of Alex's conversational move in (53) might turn out to be a good argument for veridicity, I find the theoretical lessons of these examples to be rather less clear than vFG indicate. Three objections to their interpretation stand out. First, a full evaluation of these examples would require an elaborate theory of the pragmatics of retraction and epistemic blame in conversation — but vFG do not give such a theory. Second, even in our current fragmentary state of understanding, theorists who reject the claim that epistemic statements have fixed (or indeed any) truth values have already developed plausible accounts of retraction that can explain the oddity of (53) without reference to the truth or falsity of the epistemic statement. Third, and most importantly, the same oddity arises if we replace *must* in (53) with epistemic expressions that are uncontroversially non-veridical.

Getting clear intuitions about the truth-values of epistemic statements is a delicate matter, and it is not clear to what extent disagreement, retraction, etc. are useful diagnostics here. (See, among many others, Hacking 1967; DeRose 1991; von Fintel and Gillies 2008; Knobe and Yalcin 2014.) Many authors are convinced that epistemic statements simply do not have determinate truth-values (e.g., Yalcin 2011), or that their truth-values are unstable and sensitive to the information of an evaluator (Egan 2007; MacFarlane 2011, 2014). So, while it could be that *It must be raining* is determinately false in (53), it is tendentious to ground an analysis of the distancing examples on the assumption that our intuitions in this case are tracking truth *per se*. An advocate of an expressivist or relativist semantics would presumably conclude from these examples that the felicity of distancing depends on a variety of pragmatic factors, which may or may not even make reference to truth.

For example, Yalcin (2011: Sect. 4) suggests that intuitions about disagreement involving epistemic statements may be explicable in terms of two separate concepts: whether a statement is *reasonable* in light of a speaker's evidence, and whether it is *advisable* to adopt an informational state that conforms with the content of the epistemic statement. Relatedly, Knobe and Yalcin (2014: Sect. 6) propose to account for intuitions about appropriate retraction not in terms of the truth of the original statement, but in terms of whether it would be reasonable for a speaker to disown the conversational update that the statement would have effected if accepted. On such an account, once it has been observed that it is not raining, it would be inadvisable to adopt an epistemic state that accepts either of Alex's initial statements in (53) and (54). Still, the examples differ sharply in terms of reasonableness. Alex's information at the time of utterance did support the weaker *might*-statement, but did not support the stronger *must*-statement. (After all, Alex said "It must be raining" without even taking the trivial step of looking out the window.) If we suppose that such pragmatic factors are implicated in our intuitions about the felicity of Alex's replies—rather than truth *per se*—these examples may be explicable without even supposing that the statements admit of truth and falsity, as Yalcin would have it.

Even a contextualist could adopt this style of analysis, supposing that these statements do admit of truth, but that our intuitions about appropriate distancing are not sensitive to truth alone, or at all. In fact this seems like the right move on independent grounds, since there is a direct argument showing that the felicity of distancing does not speak to truth: the relevant examples remain very odd when a clearly non-veridical certainty expression is used, such as *I am absolutely certain*.

(56) a. *Alex*: I am absolutely certain that it's raining.

b. *Billy*: [opens curtains] No it isn't. You were wrong.

c. *Alex*: ?? I was not! Look, I didn't say it *was* raining. I only said I was *absolutely certain* that it was. Stop picking on me!

Alex's response in this dialogue is at least as bizarre as her response in (54). But we can't pin the oddity of this dialogue on Alex's having said something *false*. What she said may well have been true, but Billy's criticism sticks: Alex's strong subjective certainty was not reasonable or evidentially well-grounded. This point lends further support to the suggestion above that intuitions about such dialogues do not target (just) utterance truth, but some possibly complex constellation of pragmatic factors that are at most correlated with truth.

While I do not know what exactly is going on in these dialogues, a complex variety of social, linguistic, and epistemological variables appear to be implicated in intuitions about the appropriateness of distancing. There is every reason to expect that a full theory of the conversational pragmatics of epistemic statements will be able to deal with these examples without referring exclusively to the truth-values of the initial epistemic statements. Perhaps, as (56) hints, it will not make direct reference to their truth at all.

5 Some existing accounts

5.1 Von Fintel and Gillies 2010

As noted above, several earlier authors indicate a connection between *must* and reasoning from indirect evidence (Karttunen 1972; Palmer 1979). However, vFG seem to be the first to argue explicitly that the *only* "weakening" factor is indirectness. To reinforce this claim, vFG propose a positive account on which *must* is a universal quantifier over a domain of worlds $B_{\mathcal{K}}$, while also encoding a presupposition of indirect evidence. The primary attraction of the theory is that it encodes *must*'s requirement of indirect evidence in a way that is consistent with vFG's claim that *must p* entails that p is known (and so true and certain).

In vFG's (first) account, there is a set of propositions $\mathcal{K}$, the *kernel*, each of which is known directly (or "directly enough"). They assume that the modal base $B_{\mathcal{K}}$ is the deductive closure of $\mathcal{K}$, i.e., $B_{\mathcal{K}} = \bigcap \mathcal{K}$. This assumption makes sense if we allow that whatever is known is known either directly or by deduction from something known directly (enough).

vFG then define truth-conditions for a simple epistemic *must* sentence:

(57) $$\textit{must } p \begin{cases} \text{is undefined if } \mathcal{K} \text{ directly settles } p; \\ \text{otherwise, is true iff } p \supseteq B_{\mathcal{K}} \text{ (i.e., if } B_{\mathcal{K}} \text{ entails } p). \end{cases}$$

This definition captures indirectness via the presupposition that p is not settled directly by $\mathcal{K}$. When this presupposition is satisfied, it asserts that p is known. What precisely this comes down to depends on how we define "directly settled". vFG define this notion in terms of entailment by, or contradiction with, some proposition in $\mathcal{K}$.

(58) $\mathcal{K}$ directly settles p iff there is some $q \in \mathcal{K}$ such that either (a) $q \subseteq p$, or (b) $q \cap p = \varnothing$.

To illustrate the application of this semantics to a concrete case, vFG introduce Billy, who has just seen people coming into her office with wet raincoats and muddy boots. She utters "It must be raining." The semantics predicts that this utterance is a presupposition failure if Billy has directly observed that it is raining, or that it is not raining. Otherwise, it is true just in case Billy has access to some kernel $\mathcal{K}$ of true and directly observed propositions such that $\bigcap \mathcal{K}$ entails that it's raining. Specifically, vFG suggest that Billy's utterance is appropriate and true if her kernel is as follows:

(59) $\mathcal{K}_{\text{Billy}} = \{$*people are coming in with wet raincoats and muddy boots*, *people only come in with wet raincoats and muddy boots when it's raining*$\}$

The deductive closure of any such $\mathcal{K}$ must also include the proposition *It's raining*, and so vFG predict that *It must be raining* is true. In addition, if $\mathcal{K}_{\text{Billy}}$ accurately represents Billy's epistemic state, "I know that it is raining" would be equally true in Billy's mouth. Assuming that known propositions are believed with certainty, Billy must be certain of *It is raining*. Finally, since Billy can only know true things, the truth of *It must be raining* entails that it is in fact raining. This semantics thus settles the *Knowledge*, *Veridicity*, *Certainty*, and *Impossibility* questions from Sect. 1 in the affirmative.

vFG also give a second, more complex account of what it means for a p to be directly settled. In this implementation, the kernel is a set of issues (polar questions) that have been directly settled, and the presupposition of *must p* is that the question whether p is not among these issues. While the second implementation differs in some respects from the first (vFG pp. 379-380), these differences do not make a difference for anything that follows, as far as I can see. So, I will concentrate for the moment on the simpler account in (57) and (58). (The core idea of the second will, however, re-appear in Sect. 6 as a high-level inspiration for my own proposal.)

Under either implementation, this semantics for *must* cannot be quite right, for now-familiar reasons. The corpus evidence given above shows that *must p* is compatible with an implicit or explicit indication of lack of knowledge of p, of non-maximal confidence in p, and of the possiblity of p. Similarly, most of our experimental participants were willing to endorse *Bill must not have won* despite knowing that *Bill did not win* had non-maximal probability, and even while other participants with the same information overwhelmingly endorsed *It is possible that Bill won*. At the same time, the large majority of participants rejected *We know that Bill did not win* and *It is certain that Bill did not win* in the same scenario.

A different, perhaps more theoretically helpful gloss on this problem is that vFG's account has difficulty explaining the frequent use of *must* to report inductive inferences. In the genealogical examples (3) and (4), for instance, *must* is used to mark statements which are reasonable to believe because their truth would readily *explain* the available evidence, but which are not entailed by anything known to the author. In these examples it is difficult to even imagine a set of propositions that the author could plausibly have observed directly — or indeed "directly enough" — which would also entail the truth of the statement marked with *must*. It is also implausible to suppose that the authors have an *erroneous* belief that they have access to the required set of propositions. As we noted in Sect. 2.1, experienced genealogists are well-aware of the possibility of error in their studied inferences, but they use *must* nonetheless.

vFG's theory is too restrictive to deal with *must* in the wild. Ultimately, I will argue, we can keep the good parts of this analysis while getting the facts right by retaining the idea that *must* marks inference, but discarding the assumption that inference = deduction.

5.2 Assumptions and normalcy

A conservative modification to vFG's theory designed to deal with these problems might proceed along the following lines. We can maintain the assumption that *must p* is a presupposition failure if p or $\neg p$ is entailed by anything that has been observed directly. However, in this modification *must*'s domain of quantification is no longer the set of all epistemically possible worlds $\mathcal{E}$ — which, we may continue to assume, is just $\bigcap\mathcal{K}$, the conjunction of all the propositions known directly (enough). Instead, the domain of *must* is $\bigcap\mathcal{K}\cap\bigcap\mathcal{A}$, a subset of the epistemically possible worlds containing all those that are compatible with each of a set of assumptions $\mathcal{A}$. (We are simplifying, inessentially, by assuming that the propositions in $\mathcal{A}$ are consistent with each other and with the directly known propositions.) The assumptions do not have to be known, or even believed with maximal confidence. This gives us a straightforward account of the numerous examples given above in which *must p* was used even while it was implied or explicitly asserted that there was non-zero doubt in the truth of p.

A semantics of this form could be given several different interpretations that are nevertheless closely related formally. One is to think of the elements of $\mathcal{A}$ in terms of "assumptions", as suggested above. Here, the propositions in $\mathcal{A}$ would be ones that are assumed temporarily, but which we would be willing to suspend if they were explicitly questioned. (See Franke and de Jager 2007, 2011, de Jager 2009 for a theory along these lines, though with a somewhat different formal implementation.) Alternatively, we might think of this account an extension to epistemic modals of von Fintel's (2001) concept of a "modal horizon": some worlds, while not epistemically impossible, are so remote that they are not taken into consideration for certain modal purposes. The too-remote worlds would be those that we are assuming do not hold, without knowing that they do not: $\bigcap\mathcal{K}-\bigcap\mathcal{A}$.

A third possible interpretation is in terms of Kratzer's (1991) notion of an "ordering source". We can think of $\mathcal{K}$ as the output of the modal base function f, and $\mathcal{A}$ as the output of the ordering source function g, when each is fed the evaluation world w as an argument. On this interpretation, vFG's "kernel" is a special case of a realistic modal base, and the propositions in $\mathcal{A}$ are the relevant "norms" or "expectations". If we simplify by making the limit assumption and assuming consistency, the domain of *must* can be thought of as the set of maximally normal epistemically possible worlds.

(60) $$\textit{must}\ p \begin{cases} \text{is undefined if } \exists q \in f(w) : q \subseteq p \vee p \cap q = \varnothing \\ \text{otherwise, is true iff } p \supseteq \bigcap f(w) \cap \bigcap g(w). \end{cases}$$

$f(w)$ is the same here as vFG's kernel $\mathcal{K}$, and the indirectness requirement is implemented in the same way: no proposition in $f(w)$ can entail or contradict p. $g(w)$ has the same formal role as the set of assumptions $\mathcal{A}$ in the above. (Kratzer (1991) also offers mechanisms for dealing with inconsistent assumptions and cases where there are no maximal worlds, rendering her official proposal slightly more complex. However, this definition captures the crucial points for the comparison with vFG and others.)

A modification to vFG's account along these lines has a number of useful features. First, it demonstrates that their theory is fully modular: it is straightforward to combine a weak semantics for *must* along the lines of Kratzer 1991 with the indirectness requirement that vFG give strong motivation for.

Second, the proposal just sketched (under any interpretation of the parameters) makes room for an analysis of the examples discussed above in which *must p* is conjoined with an expression indicating that p is uncertain, or that $\neg p$ is possible. The strategy is to treat *might* as *must*'s dual, denoting an existential quantifier over $\bigcap f(w) \cap \bigcap g(w)$. *Possible*, on the other hand, picks out an existential quantifier over the entire set of epistemically possible worlds $\bigcap f(w)$. As a result, *might p* asymmetrically entails *p is possible*, and *must p but it is possible that not-p* is consistent. In addition, on the assumption that *certain* is the dual of *possible*, this account would also leave room for examples of the form *must p but I am uncertain (of p)*, such as (36).[9]

This interpretation of *must* also does a reasonable job of accounting for this item's use to mark inductive inferences. For example, I noted above that authors frequently use *must p* despite clearly not being certain of p's truth. Serious genealogists, for example, would have little difficulty in listing off a variety of remote possibilities which if true would invalidate a conclusion that they have just marked with *must*. I argued above that the genealogical examples that we have seen are better understood as conveying that a conclusion is (a) strong, (b) the strongest available, and (c) the result of indirect inference.

A weak semantics that treats *must* as a quantifier over a subset of $\mathcal{E}$ is able to implement this analysis to a considerable extent. The conjunction of the strongest directly known proposition ($\bigcap f(w)$) with the strongest reasonable assumption ($\bigcap g(w)$) generates a set that can be thought of as the strongest reasonable conclusion that we can draw. In other words, on this interpretation anything that is entailed by $\bigcap f(w) \cap \bigcap g(w)$ would count as a "reasonable inference". Furthermore, this conclusion is the strongest conclusion available (desideratum (b)). In at least some cases this feature might explain why *must p* often conveys that p is the best explanation of the available evidence. In addition, if we know that p is entailed by this set because *must p* was appropriately asserted, then the reasonable conclusion that p must also be justified in some way other than through direct observation (desideratum (c)). Otherwise, there would be a presupposition failure.

5.3 The high probability requirement

What is still missing in the account just sketched is desideratum (a), the requirement that p be believed with high confidence before *must p* is appropriate. Unless we impose some additional constraints, nothing in this proposal prevents *must p* from being true even while p is probably, or even almost certainly, false. This is surely not appropriate, though: *must p* cannot be true when p is probably false. Even though the essentially Kratzerian account just sketched does a good job of

9 Note, however, that it is not really plausible to treat *certain* as a universal quantifier: as Lassiter (2010) and Klecha (2014) discuss, *certain* is a scalar adjective which maps propositions to a scale of certainty. However, if we adopt the additional probabilistic constraints motivated below, the account just described ports over directly to a scale-based treatment—e.g., any of the three accounts discussed by Lassiter (2016): ch.5.

combining semantic weakness with the requirement of indirect evidence, it conflicts with a very basic intuition about the meaning of *must*. If something *must* be the case, then it is highly probable, and so (as a consequence) more likely than not. (See Lassiter 2015 for a detailed discussion of this point and its consequences.)

A first attempt to repair the situation might be to require that the assumptions be highly probable ($\forall a \in g(w) : P(a) \geq \theta$, where θ is some high probability threshold). However, this is not sufficient unless we also impose constraints on the number of assumptions that can be made. No matter how high θ is, if it is less than 1 then it will be possible to construct a set of consistent assumptions, all of which have probability $\geq \theta$, but whose conjunction has very low probability. For example, if there are 1000 raffle tickets $t_1, \ldots, t_{1000}$ and $\theta = .995$, then each member of the set $g(w) = \{t_1 \text{ did not win}, t_2 \text{ did not win}, \ldots, t_{999} \text{ did not win}\}$ will have sufficiently high probability. If this is a permissible set of assumptions, then *Ticket* t_{1000} *must have won* could come out true because we are permitted to assume, for each other ticket, that it did not win. But this is absurd: the probability of $\bigcap g(w)$ — the probability that t_{1000} won — is just .001.

This puzzle suggests that the additional requirement that we impose cannot simply refer to the probability of individual assumptions/ordering source propositions. A slightly stronger option is to constrain the probability of the conjunction of the assumptions. Continuing to co-opt vFG's unsettledness presupposition, the result would be something like:

(61) $$\textit{must } p \begin{cases} \text{is undefined if } \exists q \in f(w) : q \subseteq p \vee p \cap q = \varnothing \\ \text{else, is true iff } (p \supseteq \bigcap f(w) \cap \bigcap g(w)) \wedge P(\bigcap g(w)) \geq \theta, \end{cases}$$

where θ is a high probability threshold.

(We continue to simplify here by ignoring the possibility of inconsistent or infinite premise sets.)

The definition in (61) still isn't quite right, though. It predicts that *must p* will be false if we are assuming too much — that is, if $P(\bigcap g(w))$ is too low — even in cases where many of the assumptions are irrelevant to the truth of p. But surely we don't want *must p* to come out false simply because of irrelevant features of the contextual parameters.

A better way to capture the intent of this proposal is to adapt Swanson's (2015) semantics for epistemic *must*. Rather than constraining the probability of $g(w)$, Swanson suggests simply adding to Kratzer's (1991) semantics an additional requirement that the prejacent have very high probability. Roughly:

(62) $$\textit{must } p \begin{cases} \text{is undefined if } \exists q \in f(w) : q \subseteq p \vee p \cap q = \varnothing \\ \text{else, is true iff } (p \supseteq \bigcap f(w) \cap \bigcap g(w)) \wedge P(p) \geq \theta, \end{cases}$$

where θ is a high probability threshold.

The proposal in (62) maintains several attractive features of the assumption-based solution. For example, it does not require probability 1, and so it leaves room for *must p* to be true when p does not follow from the directly known propositions, as long as it is very probable and also follows from the known propositions together with the assumptions.

When we turn to *must*'s dual *might*, however, (62) encounters some difficulty. Recall that Sects. 2 and 3 gave evidence that *might* is stronger than *possible*. Given that *might* is *must*'s dual, the conjunctive definition in (62) predicts a disjunctive meaning for *might q*: either the first conjunct of the *must* definition fails to hold of $\neg q$, or the second does. To reduce clutter, let $\mathcal{D}$ represent $\bigcap f(w) \cap \bigcap g(w)$, the set of worlds compatible with what is known *and* what is assumed. Also, let's assume that the presupposition is satisfied: q is not directly settled by $f(w)$. We then have the equivalences:

(63)

$$\begin{aligned} \textit{might } q &\Leftrightarrow \neg \textit{must} \neg q \\ &\Leftrightarrow \neg [\neg q \supseteq \mathcal{D} \wedge P(\neg q) \geq \theta] \\ &\Leftrightarrow [q \cap \mathcal{D} \neq \varnothing] \vee [P(q) > (1-\theta)] \end{aligned}$$

The prediction is that, whenever its indirectness presupposition is satisfied, *might q* is true if either of the following conditions holds: either

(64) a. q is compatible with everything that is known and everything that is assumed, or
b. q is probable enough that its negation would not meet the probabilistic threshold associated with *must*.

The proposal in (62) plus duality thus predicts that *might q* is true whenever q is compatible with the relevant knowledge and assumptions, even if the probability of q is zero. In addition, *might q* can be true even when it is being assumed that q is false, as long as it has significant probability. It is hard to be sure what the practical import of these features is, since we have no theory of how probability assignments constrain assumptions/ordering sources. However, the disjunctive interpretation of *might* does have the straightforward prediction that *might* cannot be stronger than *possible*, under any of the usual interpretations of the latter item. For example, each of the following has been suggested as an interpretation of *q is possible*.

(65) a. $q \cap \bigcap f(w) \neq \varnothing$ (von Fintel and Gillies 2010, cf. Yalcin 2010)
b. $q \cap \bigcap f(w) \cap \bigcap g(w) \neq \varnothing$ (inspired by Kratzer 1991)
c. $P(q) > 0$ (Yalcin 2005; Lassiter 2010)

The dual of Swanson's *must* is not strictly stronger than any of these, as (63) shows. It is logically independent of (65a) and (65c), and strictly weaker than (65b).

One possible response is to deny duality. For example, Swanson (2015) suggests treating *might* as Kratzer's *might* conjoined with a condition of sufficiently high probability — in effect, the third line of (63) with the "$\vee$" changed to "$\wedge$". Bracketing the evidential presupposition, this gives us[10]:

(66) *might p* is true iff $[q \cap \mathcal{D} \neq \varnothing] \wedge [P(q) > (1-\theta)]$.

10 We are also bracketing some thorny but orthogonal issues about the truth-aptness of probability talk. Swanson's (2015) discussion is really focused on these issues, rather than the specific lexical semantic proposals that we are discussing here: he argues for a constraint-based rather than truth-conditional approach to probabilistic language generally. The entries given in the main text are couched in the language of truth-conditions, but they could—as far as I can see—be readily adapted to Swanson's broader framework, and nothing I say here should be construed as an argument against the latter.

This is a possible solution, but I am hesitant to abandon the duality of *must* and *might* so quickly. In addition, there is a potential empirical problem: the proposal predicts that it should be possible for both *must p* and *might not-p* to be false. Consider a situation in which it follows from $\mathcal{D}$, the directly known propositions and the relevant assumptions, that Bill is at home—but this proposition is not probable enough to meet the threshold θ. (This could, but need not, happen if the conjunction of our assumptions has probability lower than θ.)

Under these conditions, *Bill must be at home* will come out false by (62), even though it is entailed by $\mathcal{D}$, because the probability of *Bill is at home* does not meet or exceed θ. So far so good. However, *Bill might not be at home* also comes out as false by (66). This is because, even though *Bill is not at home* has probability greater than θ, it is incompatible with $\mathcal{D}$. Intuitively, though, it seems absurd to deny both *Bill must be at home* and *Bill might not be at home*.

While we should not place too much weight on such intuitions, they provide at least a *prima facie* reason to seek out an alternative. The next section will sketch one way to rescue duality while maintaining one of the most attractive features of Swanson's (2015) brief suggestions about the lexical semantics of *must* and *might*: his way of encoding the connection between *must* and high probability. This approach also makes it possible to incorporate vFG's treatment of *must*'s requirement of indirect evidence.

6 Threshold semantics in structured probabilistic models

Our goal is to find a semantics on which (a) *must* requires high but non-maximal probability, (b) *must* is associated with an indirectness requirement, and (c) the predicted dual is a plausible interpretation of *might*, which is stronger than some plausible interpretation of *possible*. In fact, an earlier proposal by Swanson (2006) has two of these features — (a) and (c). Swanson's simple proposal was to identify *must p* with the condition that the probability of *p* is sufficiently high.

(67) *must p* is true iff $P(p) \geq \theta$, where θ is a high probability threshold.

By duality, this definition predicts a similarly straightforward interpretation of *might*:

(68) *might p* is true iff $P(p) > (1 - \theta)$, where θ is a high probability threshold.

And the latter is stronger than at least one reasonable interpretation of *possible* (Yalcin 2005; Lassiter 2010).

(69) *p is possible* is true iff $P(p) > 0$.

This section sketches out a way to extend this idea to include the indirectness requirement using probabilistic graphical models, a format for knowledge representation that is widely employed in artificial intelligence and cognitive science. This approach has the benefits of the accounts due to vFG, Kratzer, and Swanson discussed above, but it also avoids the problems that we noted for these.

6.1 Question-based probabilistic information dynamics

Probabilistic graphical models (PGMs) are a standard tool for knowledge representation and inference in artificial intelligence and computational psychology.[11] In this approach, as in other Bayesian formalisms, information states are represented using a set of epistemically possible worlds $\mathcal{E}$ together with a measure function P which (a) takes propositions to $[0,1]$, (b) assigns 1 to $\mathcal{E}$, and (c) is additive—$P(A \cup B) = P(A) + P(B)$ if $A \cap B = \varnothing$. Update via learning is modeled using conditionalization. If the probability distribution is $P(\cdot)$ and then B is learned ($P(B) \neq 0$), the distribution is updated to incorporate the information that B is true. Afterwards, for any proposition C, $P(C \mid B) = P(B \wedge C)/P(B)$.

PGMs add structure to this basic picture in two ways. First, they contain a set of "variables", each of which partitions $\mathcal{E}$. Formally, a variable is simply a Groenendijk and Stokhof 1984 question denotation restricted to $\mathcal{E}$ (see van Rooij 2003).

(70) **Definition** (Variable/Question). A variable Q is a set of propositions ("answers") where
  a. all answers in Q are subsets of $\mathcal{E}$;
  b. any two distinct answers to Q are mutually exclusive;
  c. Q covers $\mathcal{E}$: one of the answers to Q is true.

Second, PGMs add "arrows" representing dependencies between variables/questions. If there is an arrow from Q_1 to Q_2, then Q_1 and Q_2 are dependent in the sense that learning the value of one can influence the probability distribution on values of the other. Independent variables have the converse property: if Q_1 and Q_2 are independent then there is no answer to Q_1 which, if it were learned, would lead the learner to modify the probability that she assigns to any answer to Q_2 (and vice versa). (In)dependence assumptions will not play a major role here, but they are critical for many applications of PGMs.

My proposal builds on the PGM framework, while also taking inspiration from vFG's second, issue-based theory of *must* as well as recent formal models of awareness (Franke and de Jager 2007; de Jager 2009; Yalcin 2011). I assume that agents represent only some of the possible questions—that is, that they are sensitive only to certain issues at a given time. Let the relevant set of questions be $\mathcal{V}$. As a set of question meanings, $\mathcal{V}$ is also a set of variables in the PGM sense. I assume further that $\mathcal{V}$ is partitioned into two sets: a set $\mathcal{V}_{\mathcal{D}}$, the variables whose true values have been observed directly, and $\mathcal{V}_{\mathcal{I}} = \mathcal{V} - \mathcal{V}_{\mathcal{D}}$, whose values have not been observed directly and must be inferred.

For any m, let P_m denote the relevant probability measure at time t_m. Suppose that, at t_n, the answer to question $\mathcal{Q}$ is directly observed (and no other new information is acquired). The following conditions must then hold:

- If $\mathcal{Q}$ is in $\mathcal{V}_{\mathcal{I}}$ at t_n, it is in $\mathcal{V}_{\mathcal{D}}$ at t_k for $k > n$.

11 AI: Pearl 1988, 2000; Spirtes, Glymour and Scheines 1993; Koller and Friedman 2009; Russell and Norvig 2010, and many more. Psychology: Glymour 2001; Sloman 2005; Gopnik and Schultz 2007; Danks 2014, and—with additional structure—Tenenbaum et al. 2011. The type of probabilistic graphical models that I discuss here are also known as "Bayesian networks". They are not the only kind of PGM, but they are one of the most prominent and intuitive.

- There is some answer $q \in \mathcal{Q}$ such that $P_k(q) = 1$ for all $k > n$.
- Suppose, after t_n, that no other direct observation is made until t_p. Then, for all t_k such that $t_n < t_k < t_p$, and for all $\mathcal{Q}' \in \mathcal{V}$, $P_k(\mathcal{Q}') = P_n(\mathcal{Q}' \mid \mathcal{Q})$.

The first condition updates the direct/indirect division: when an agent observes the answer to a question about which she has previously made only probabilistic inferences, that question is registered thereafter as one whose answer has been directly observed. The second requires that direct observation is uncertainty-free and that observed answers are not forgotten. (These are helpful but not insignificant simplifications). For instance, if we are representing the question *Is it raining at t_m?* and then observe that it's raining at t_m, then $P(rain-at-t_m) = 1$ for all times thereafter. The third condition requires that information dynamics proceed exclusively by Bayesian update on direct observations, ruling out non-Bayesian update and clairvoyance.

A further notable feature of the third condition is that, at any time t_m, the distribution P_m is equal to the initial distribution P updated by the values of all directly observed variables: $P_m(\mathcal{Q}) = P(\mathcal{Q} \mid \mathcal{V}^m_{\mathcal{D}})$, where $\mathcal{V}^m_{\mathcal{D}}$ is the set of all questions whose values have been observed up to time t_m. Given this, we can simply write $P(\cdot \mid \mathcal{V}_{\mathcal{D}})$ for the relevant probability distribution at any time, letting the choice of $\mathcal{V}_{\mathcal{D}}$ take care of the temporal bookkeeping.

6.2 Capturing *must*

Using the PGM-based model of information dynamics, it is straightforward to modify Swanson's (2006) scalar semantics for *must* to meet our desiderata.

I assume that assertions are required to address a Question Under Discussion (QUD) which is determined by the discourse context, as proposed by Ginzburg (1995a,b), van Kuppevelt (1995), Roberts (2012), and Beaver and Clark (2008). An assertion of *must p* presupposes that p addresses the QUD—simplifying, that it is an element of (direct answer to) the QUD. This is added to *must*'s special presupposition to the effect that the QUD is not one whose answer has been directly observed.

(71) $$\textit{must } p \begin{cases} \text{is a presupposition failure if } p \notin QUD \text{ or } QUD \in \mathcal{V}_{\mathcal{D}}; \\ \text{otherwise, is true iff } P(p \mid \mathcal{V}_{\mathcal{D}}) \geq \theta, \end{cases}$$

where θ is a high probability threshold.

The presupposition is an elaboration of vFG's condition, in their second implementation of the kernel-based semantics, that the question whether p not be settled by anything that has been directly observed. The truth-condition is effectively Swanson's (Swanson 2006), except that it explicitly enforces the condition that the relevant information state is the one derived from the initial state $P(\cdot)$ by conditioning on the direct observations up to the relevant time.

This proposal predicts the core intuitions driving vFG's account. For example, consider again the scenario in which Billy has directly observed that it is raining. As vFG point out, she cannot say *It must be raining*. On the present account, the reason is that the QUD is *Is it raining?*, and the fact that she has observed rain directly ensures that the question *Is it raining?* is an element of $\mathcal{V}_{\mathcal{D}}$. As a result, her utterance triggers a presupposition failure.

On the other hand, consider the situation in which Billy has observed people coming into the room with wet raincoats and rain boots. If she has not directly observed rain or lack thereof, the question *Is it raining?* is safely in $\mathcal{V}_{\mathcal{I}}$, and the presupposition of *It must be raining* can be satisfied. Given this, the question that remains to answer is, in effect, whether *It is raining* has sufficiently high probability conditional on $\mathcal{V}_{\mathcal{D}}$, which represents the sum total of Billy's direct experience of the world. If it does, then *It must be raining* is true. This truth-condition thus allows for the observation from above that *must* can be used to mark inductive inference—here, that *It is raining* is a reasonable conclusion for Billy to draw, though defeasibly, in light of her evidence.

On this semantics, *must* is upward monotone (as long as all relevant presuppositions are satisfied). It is a theorem of the probability calculus that $P(p) \leq P(p \vee q)$. So, as long as θ is held fixed, *must p* entails *must* $(p \vee q)$. This much is shared with most theories, including both von Fintel and Gillies 2010 and Kratzer 1991. However, this *must* has at least one unusual property: it is not closed under conjunction. That is, the following argument is not logically valid.

(72) a. Must p.
b. Must q.
c. $\therefore$ Must $p \wedge q$.

For example, if $\theta = .99$, and p and q are probabilistically independent and both have probability .99, then *must p* and *must q* are both true as long as the indirectness presupposition is satisfied. In this case, $p \wedge q$ will have probability $.99 \times .99 = .9801$. This is less than θ, and so *must* $p \wedge q$ is false.

While this might seem to be a problem on intuitive grounds, we have already seen reason not to value intuitions too much where such subtle empirical questions are involved (compare Sect. 4.2 above). In addition, the predictions about closure are not as neat as (72) might suggest. The assumption that θ has a single, known value surely will not characterize any real communicative interaction. Rather, there will generally be uncertainty about the value of θ, and interpreters must integrate this uncertainty into their judgments in some way. In effect, the theory given here treats *must* as a vague scalar expression. Given this, we expect judgments to display the usual sorites-like behavior: small changes in probability, such as the difference between .99 and .9801, will not lead to a drastic change in the intuitive acceptability of a *must* claim unless there is an extraordinary level of precision in the conversation about the values of free parameters. But when there is significant uncertainty about θ, the fact that it is clearly at or below .99 (so that the premises are true) means that there will also be reasonable candidate values of θ at or below any value just slightly below .99, such as .9801. It is only when the decrease in probability is quite large that we would expect a precipitous decline in acceptability. In fact, a large decline in probability between the premises and conclusion of (72) is impossible unless θ is low, given that only two propositions are being aggregated within the scope of *must* (compare Adams 1975).[12] However, if many such conjuncts were being aggregrated under *must*—say, the 1,000 propositions of the form *x did not win* for each ticket-holder *x*—their conjunction could have low or even zero probability, consistent with each having the probability required for *must* to hold. See Edgington 1997 and Lassiter and Goodman

12 Specifically, for any choice of p and q where $P(p) \geq \theta$ and $P(q) \geq \theta$, we have $P(p \wedge q) \geq 1 - 2 \times (1 - \theta)$. In other words, when p and q both have probability within $1 - \theta$ of the maximal value 1, the gap between 1 and their conjunction's probability cannot be more than twice as great.

2015a for a development of closely related points involving non-modal scalar expressions, and Sect. 7.3 below for some further relevant discussion.[13]

Assuming duality of *must* and *might*, (71) predicts a simple, plausible interpretation of the latter. This interpretation maintains the indirectness presupposition and associates *must p* with a requirement that p has probability greater than $1-\theta$.

(73) *might p* $\begin{cases} \text{is a presupposition failure if } p \notin \mathcal{Q} \text{ or } \mathcal{Q} \in \mathcal{V}_{\mathcal{D}}; \\ \text{otherwise, is true iff } P(p \mid \mathcal{V}_{\mathcal{D}}) > (1-\theta), \end{cases}$

where θ is a high probability threshold.

This is at least as strong as *possible* on a lightly modified version of the interpretation suggested by Yalcin (2005) and Lassiter (2010):

(74) *p is possible* is true iff $P(p \mid \mathcal{V}_{\mathcal{D}}) > 0$.

If $\theta = 1$, *might* requires probability greater than 0, and *possible* and *might* have the same truth-condition. When $\theta < 1$ — presumably, the normal case — *might* is strictly stronger than *possible*, since it requires a probability greater than some value $1-\theta$ which exceeds 0. The same result holds if we treat *possible* differently as an existential quantifier over epistemically possible worlds. This semantics requires a slightly more elaborate representaton (see, e.g., Yalcin 2010 for some relevant discussion), but it gives *possible* a strictly weaker meaning than (74). So, this theory continues to enforce the empirically motivated condition that *might* is stronger than *possible*.

7 Some loose ends

7.1 Veridicity and assertion

So far, we have avoided the crucial issue of whether *must* is veridical. I argued above that the empirical evidence from our corpus and experimental investigations did not speak to the issue, and that the arguments given by vFG in favor of veridicity were not decisive. In this section I will consider the consequences of adding to this semantics a requirement that *must p* implies *p*. The upshot is that, in light of the conclusions we have reached, it is technically feasible to treat *must* as veridical. However, doing so appears to require significant commitments with respect to the theory of assertion that some may consider undesirable.

13 At first glance, the lottery experiment reported in Sect. 3 seems to provide some evidence for failure of closure. A majority of the participants endorsed the prompt *Bill must not have won* when Bill had one of 1000 tickets, so that his probability of not winning was 999/1000 = .999. Since there was nothing special about Bill or his ticket in this scenario, presumably the same participants would, if asked, have been willing to endorse *Judy must not have won, Sam must not have won*, and so on for the other 999 individuals, each of whom had probability .999 of not winning. Still, these participants would presumably not have been willing to endorse *It must be that no one won*. However, the force of this point is weakened considerably by the fact that the prompt *Bill did not win* had a similar, fairly high rate of acceptance as the *must not* prompt. Surely we would not want to conclude from this observation, together with the clear truth of *Someone won*, that truth itself is not closed under conjunction. We should presumably look to participants' interpretation of "accept" for clues: acceptance is tracking something slightly weaker than truth. The experimental results do not strongly favor the failure of closure under conjunction, then, but they are certainly consistent with this property of the theory of *must* given here.

We can make *must* veridical by adding p as a conjunct to the definition in (71).

(75) $$\textit{must } p \begin{cases} \text{is a presupposition failure if } p \notin QUD \text{ or } QUD \in \mathcal{V}_{\mathcal{D}}; \\ \text{otherwise, is true iff } p \wedge P(p \mid \mathcal{V}_{\mathcal{D}}) \geq \theta, \end{cases}$$

where θ is a high probability threshold.

By duality, we then predict that *might* should be associated with the same presupposition and the assertive content in (76).

(76)

$$\begin{aligned} \textit{might } p &\Leftrightarrow \neg \textit{must} \neg p \\ &\Leftrightarrow \neg[\neg p \wedge P(\neg p \mid \mathcal{V}_{\mathcal{D}}) \geq \theta] \\ &\Leftrightarrow p \vee [P(p \mid \mathcal{V}_{\mathcal{D}}) > (1-\theta)] \end{aligned}$$

The prediction is that *might p* is true whenever p is—and, in addition, when p is false but it has sufficient probability according to the relevant probability measure. The prediction that p implies *might p* is in agreement with the classical semantics for *might* that vFG endorse, according to which this item is an existential quantifier over a set of worlds which includes the actual world. On both accounts, if you judge *It might be raining* false—but it is in fact raining—you are simply wrong about what might be the case.

When we turn from semantics to the pragmatics of assertion, a veridical scalar theory of *must* makes interesting predictions. A popular (though not uncontroversial) theory of assertion holds that one should assert p only if one knows that p is true (Williamson 2000; Hawthorne 2004; Turri 2016). If this is right, then one should assert *must p* only if one knows that *must p* is true. Now, (75) holds that *must p* entails p. I assume that knowledge is closed under known entailment: if x knows p and x knows that p entails q, then x knows q. So, the combination of the knowledge norm with a veridical *must* predicts that one should assert *must p* only if one knows p (inter alia).

This conclusion is problematic in light of the corpus evidence reviewed in Sect. 2. There, we saw many examples indicating that speakers do use *must p* when they do not know p, but merely wish to indicate that p is highly likely in light of available evidence.

(77) I've seen these before and <u>they must be aftermarket</u> but <u>I don't know for sure</u>.

(78) <u>Must be an old DTS diesel setup</u> but <u>Im not certain</u>. Why would they have added this extra injector?

(79) <u>It's not impossible</u>, granted, but IMO <u>it is highly unlikely</u>. There <u>*must*</u> be some other cause.

If *must* is veridical and knowledge is the norm of assertion, these speakers (and those who generated many of the other examples in Sect. 2) are violating an important conversational norm. Not only are their *must*-statements in error, but they are aware that these statements are in error, and say so explicitly. Their statements should be at least as strange as modified examples with bare assertions replacing *must*.

(80) # I've seen these before and <u>they are aftermarket</u> but <u>I don't know for sure</u>.

(81) # It is an old DTS diesel setup but Im not certain.

(82) # It's not impossible, granted, but IMO it is highly unlikely. There is some other cause.

But these examples are all exceedingly strange, indeed Moore's-paradoxical.

If we find it unacceptable to conclude that ordinary speakers are systematically in error—and we should—there are two obvious options: we can reject the knowledge norm or the validity of the inference from *must p* to *p* (or indeed both).

If we reject veridicity of *must*, we have to deal with the arguments that vFG give in its favor. I argued in Sects. 4.3-4.4 that both of vFG's main arguments along these lines rely on unstated and potentially problematic assumptions. If so, it is possible to account for the relevant intuitions without assuming veridicity. While I don't doubt that there is much more to be said here, denying veridicity seems like a promising theoretical option, especially for a theorist who wishes to account for the empirical facts in a way that is consistent with the knowledge norm.

If we wish to maintain veridicity, we should look for an alternative to the knowledge norm of assertion. There are many accounts on the market. A certainty-based theory of assertion (Stanley 2008; Beddor 2015) would not fare better than the knowledge account, given evidence above indicating that speakers are willing to assert *must p* in the absence of certainty (Sect. 2). However, there are various weaker theories that connect assertibility with other factors: rational credibility (Lewis 1979; Douven 2006), reasonable belief (Lackey 2007), or simple truth (Weiner 2005). If one of these theories of assertion turns out to be correct, then it may be possible to analyze *must* as in (76) while also making sense of examples of the form *must p, but I don't know that for sure*; *must p, but it's uncertain*; and so on. Here is a brief sketch of how this might go for one of these accounts, chosen more or less at random.

In general, we expect that someone who is striving to conform her behavior to norm N will perform action a only if she is confident enough, for practical purposes, that a conforms to N. In the case of assertion, this means that a speaker who is striving to conform to a given norm of assertion will not utter u unless she is confident enough, for practical purposes, that the relevant norm(s) are satisfied in the context in which u is to be uttered (compare Williamson 2000: 245-6). In this light, suppose that the truth norm is correct. If (75) is the right meaning for *must*, we expect that a speaker who is trying to conform to this norm will not assert *must p* unless she is sufficiently confident that (a) its indirectness presupposition is satisfied, (b) p has high probability, and (c) p is true. The first two conditions involve our speaker's access to features of the relevant evidential state. These are not trivial—especially if the evidence is not her own—but let us suppose that they are satisfied. (For example, let the information state be the speaker's own, and assume that she has good introspective access to properties of her first-order information.) The remaining condition is that the speaker assign sufficiently high probability to p. This pragmatic condition on proper assertion of *must p* is strikingly close to one part of the truth-condition of *must p*: the requirement that $P(p) \geq \theta$, for some high θ.

In other words, if we assume the truth norm of assertion and a veridical *must*, we predict that speakers may sometimes use *must p* when they assign very high, but perhaps non-maximal, probability to p. This condition would exist alongside the very similar high-probability requirement in the semantics of *must* that I have proposed, and the two will no doubt be difficult to disentangle empirically. Similar predictions could, I think, be extracted from several of the other theories of assertion in the literature. I do not want to pursue issues involving assertion further here. The key point is that, if we adopt the veridical theory of *must* in (75) along with a theory of assertion on which appropriate assertion does not require knowledge or certainty, we may be able to make sense of the empirical evidence in Sect. 2. These speakers' productions are reasonable because their assertions do not strive for knowledge or certainty.

A decision on which of these analytic routes is best will have to await further work. For our purposes, it is enough to note that the empirical evidence reviewed in this paper here generates an interesting interaction between the lexical semantics of *must* and theories of assertion. One can maintain that *must* is veridical at the cost of rejecting knowledge- and certainty-based theories of assertion. Or, one can maintain these theories of assertion at the cost of denying the veridicity of *must*. Both options may represent fruitful directions for further investigation.

7.2 The status of the knowledge⇝certainty inference

In introducing the knowledge⇝certainty inference in Sect. 1, we noted briefly that the validity of this inference has been denied in some previous work. This section returns to these arguments.

Consider (6) from Unger 1971, repeated here as (83).

(83) # He knows it's raining, but he isn't certain of it.

Examples of this form have sometimes been claimed to be acceptable, in special contexts. For example, Radford (1966) suggests that we might describe a stressed-out student in an oral exam, who is having trouble bringing to mind all of her relevant knowledge, as knowing an answer despite not believing that it is the answer, and despite being uncertain at the moment of production. Still, as Radford hints (p.5) it is crucial to the intuitive force of the example that the student has learned the relevant facts effectively, and would be able to recall them, with an accompanying sense of belief and certainty, if nerves were not interfering with her memory. A similar but more extreme example might involve patients with a neurological condition called blindsight (Danckert and Rossetti 2005). Such patients are able to act appropriately in response to visual stimuli (e.g., in grasping, recognizing perceptual features, or even naming objects and words) in a part of their visual field in which they are, subjectively, completely blind. These patients know what they are seeing in some sense, even though they are clearly also uncertain in another sense—maximally so, since they lack any consciously accessible information about the relevant portion of their visual field that would permit them to form a sense of certainty.

Such issues of memory, attention, and awareness provide important data for pragmatic theorizing (see, for example, Franke and de Jager 2007, 2011; de Jager 2009; Yalcin 2011; Crone 2015). But the entailment data that plays a crucial role in semantic theorizing about epistemic modals typically involve a layer of psychological idealization, and it is not clear whether performance issues involving forgetfulness, schizophrenia, blindsight, self-deception, etc. should be part of the basic theory. That said, if it is possible to know something without being (momentarily) certain of it, and without an immediate sense of belief, this should be accounted for. Perhaps a more precise statement of the relevant principles should be slightly weaker: knowledge of p implies belief that p, and (ordinary) certainty that p, on the part of a subject who has full awareness of relevant issues and is not psychologically impaired or otherwise impeded. This version is still sufficient to undergird the empirical objections to the epistemic necessity account that I raised previously on the basis of examples involving *must* and (un)certainty. Those examples, where an author explicitly chooses to formulate an assertion conjoining *must p* with an expression entailing uncertainty about *p*, do not seem to involve inattentiveness, memory problems, or neurological disorders.

Stanley (2008) presents a number of arguments against the knowledge$\rightsquigarrow$certainty inference, in addition to the problematic "nervous student" example. Some of the examples that Stanley cites strike me as quite strange (e.g., *There are some things I know, which are only fairly certain to be true*, which he judges as acceptable). In addition, the claim that knowledge is possible without certainty would seem to predict that examples with the form of (8b)—*He knows it's raining, but he considers it possible that it isn't*—should be accurate descriptions of some non-pathological psychological states (Lewis 1996). Stanley (2005) attempts to defend this consequence by defining *possible* directly in terms of knowledge. However, this definition together with Stanley's (2008) denial of the knowledge$\rightsquigarrow$certainty inference would predict the satisfiability of bizarre conjunctions like (7), repeated here.

(84) # It's not possible that it's raining, but it's not certain that it isn't.

Another potential issue is that Stanley's overall account of the relationship between knowledge and certainty depends crucially on a certainty norm of assertion. As we discussed briefly in Sect. 7.1, this norm is problematized by a number of the corpus examples in Sect. 2, unless we reject the veridicity of *must*.

However, there is a real puzzle around Stanley's (2008) example *I know that Bill came to the party. In fact, I am/it is certain that he did* (p.38). Since it is usually odd to reinforce straightforward entailments in this way, Stanley concludes that knowledge does not imply certainty. Perhaps this suggestion will turn out to be correct. Even if it is, though, the prospects for the epistemic necessity theory of *must* remain dim. The semantic problem with corpus examples of the form *must p, but I'm uncertain that p* may be resolved, but the pragmatic objection discussed in Sects. 2.2-2.3 above would remain. The epistemic necessity theory predicts that these examples should be bizarre, since they assert *must p* while explicitly maintaining, and advising the listener to maintain, $\neg p$ as a live possibility. In addition, as we emphasized above, the experimental evidence in Sect. 3 directly contradicts the predictions of the epistemic necessity theory, in a way that does not require the knowledge$\rightsquigarrow$certainty inference to hold.

7.3 Brief speculations on threshold-setting

The account presented in Sect. 6.2 leaves one important issue unresolved: the way that the threshold θ is determined, relative to some context. While I cannot offer a fully explicit theory of how this happens, I think I can explain why such a theory is out of reach. The problem is essentially the same as the notoriously difficult issue of determining extensions for scalar expressions with context-sensitive meanings. Fortunately, there have been recent theoretical developments on the latter front that may be helpful in suggesting avenues for inquiry in the epistemic case.

Must is in pragmatic competition with a number of items that can be used to express grades of uncertainty and/or (in)directness of evidence. Many of these items are also vague scalar expressions with context-sensitive meanings, such as *(im)plausible*, *(un)likely*, *(un)certain*, *(im)possible*, and—I would argue—*might*.[14] It is reasonable to suppose that whatever pragmatic mechanisms go into estimating the vague and context-sensitive threshold values for scalar adjectives—*tall, short, heavy, warm, hot, full, wet*, etc.—are also recruited for the purpose of estimating the threshold value that is implicated in *must*'s meaning. A close analogue, on this way of thinking, would be the choice between *warm* and *hot* to describe temperature. Even though the interpretations of these items vary greatly as a function of context and reference class, lexical knowledge plays a partial role in fixing their thresholds: we know that the minimum temperature required for something to be *hot* is greater than that required to be warm, in any context. In other words, even though the *values* of θ_{warm} and θ_{hot} are not fixed by our linguistic knowledge, the *ordering* $\theta_{warm} < \theta_{hot}$ is.

While epistemic expressions may have vague and context-sensitive meanings, there are facts about their orderings that are similarly fixed lexically. For example, the scalar epistemic adjectives *probable* and *likely* have context-sensitive meanings that change with the distribution of probabilities among salient alternative sets (Teigen 1988; Windschitl and Wells 1998; Yalcin 2010; Lassiter 2011). In addition, a recent experimental examination of epistemic readings of six adjectives found evidence for the ordering $possible < plausible < probable < likely < certain/necessary$ (Lassiter and Goodman 2015b) . The obvious way to enforce the ordering $probable < likely$ is to make $\theta_{probable} < \theta_{likely}$ a lexical feature of these adjectives, analogous to $\theta_{warm} < \theta_{hot}$. This constraint will ensure that the two thresholds move in lockstep, regardless of the mechanisms and information sources that are recruited to resolve the interpretation further. We can constrain the interpretation of *must* similarly with a lexical ordering relative to other epistemic expressions: for example, θ_{must} is probably required to be greater than θ_{likely}. Perhaps it is also constrained to be less than $\theta_{certain}$.

While there is still much to be said about the way that speakers and listeners narrow in on a more precise interpretation of *must*, the comparison with scalar adjectives raises the hope that a theory of the threshold inference process for scalar adjectives could simply be plugged in to derive reasonable results for *must*. See Lassiter and Goodman 2013, 2015a for a recent pragmatic theory of threshold inference which is compatible with the present account.

8 Conclusion

The semantics of *must* is a rich topic with many connections to other epistemic items and to philosophical and psychological theories of knowledge and inference. In this paper we have addressed some — though certainly not all — of the important issues around its meaning and its semantic and pragmatic relationship to other epistemic items. On the basis of corpus evidence, an experiment, and analysis of a number of previous and new semantic proposals, I argued for the following conclusions.

First, corpus and experimental evidence indicates that *must* does not require that its propositional argument be known or certain, and may not even require that its argument be true. In fact, speakers sometimes use *must p* even while explicitly indicating that they do not know *p*, that *p* is uncertain, that *not-p* is possible, and so forth.

Second, the experiment in Sect. 3 showed that *must p* does not entail either *p is known* or *p*

14 Note that semantic scalarity does not imply grammatical gradability: *likely*, *certain*, etc. are both scalar and gradable, but *must* and *might* are not gradable. See Lassiter 2016: ch.1 for discussion of this distinction.

is certain. In addition, it showed that *must p* is compatible with $\neg p$ *is possible*, and replicated a previous finding showing that *possible* is weaker than *might*. The results problematize standard assumptions about the meanings of several of these expressions, but they do support standard duality assumptions: since *must* requires lower confidence than *certain*, *must*'s dual *might* requires greater confidence than *certain*'s dual *possible*.

Finally, we saw that existing theories of *must* due to von Fintel and Gillies (2010), Kratzer (1991), and Swanson (2015) contain important elements of a final account, but all of them have certain undesirable consequences or are otherwise incomplete. However, I argued that we can capture the best aspects of these accounts, while maintaining duality and the semantic relationships described above, if we model *must* using an version of Swanson's (2006) scalar account with additional scaffolding provided by probabilistic graphical models.

This conclusion also points to the importance of probability in the semantics of epistemic language more generally—a point made first by Yalcin (2005, 2007) and Swanson (2006), and reinforced in much subsequent work (e.g., Yalcin 2010, 2012a,b,c; Swanson 2010, 2015; Lassiter 2010, 2014, 2015, 2016; Moss 2013, 2015; Klecha 2012, 2014). I have argued that this connection is relevant not only to gradable adjectives like *likely* and *certain*, where the need for a scalar theory is hard to deny—but also to the epistemic auxiliaries *must* and *might*, where we can find indirect but nevertheless strong motivation for the presence of a scalar meaning component.

Acknowledgements I'm grateful to Kai von Fintel, Thony Gillies, John Turri, Daniel Goodhue, Deniz Rudin, and two *Natural Language Semantics* reviewers for extensive and helpful discussion, and also to Andy Egan, Cleo Condoravdi, Philippe Schlenker, Brett Sherman, Lelia Glass, Matthew Stone, Anastasia Giannikidou, three *SALT 24* reviewers, and audiences at *SALT 24* and the CSLI workshop *Perspectives on Modality*. Thanks also to Christine Bartels, whose stylistic and substantive suggestions led to major improvements in the final manuscript.

References

Adams, Ernest W. 1975. *The logic of conditionals: An application of probability to deductive logic*. Heidelberg: Springer.

Beaver, David and Brady Clark. 2008. *Sense and Sensitivity: How Focus Determines Meaning*. Malden, MA: Wiley-Blackwell.

Beddor, Robert. 2015. Certainty first. Ms., Rutgers University.

Crone, Philip. 2015. Asserting clarity and managing awareness. To appear in *Proceedings of Sinn und Bedeutung 2015*.

Danckert, James and Yves Rossetti. 2005. Blindsight in action: what can the different sub-types of blindsight tell us about the control of visually guided actions? *Neuroscience & Biobehavioral Reviews* 29(7). 1035–1046.

Danks, David. 2014. *Unifying the mind: Cognitive representations as graphical models*. MIT Press.

Degen, Judith. 2015. Investigating the distribution of *some* (but not *all*) implicatures using corpora and web-based methods. *Semantics and Pragmatics* 8(11). 1–55.

Degen, Judith and Michael K Tanenhaus. 2015. Availability of alternatives and the processing of scalar implicatures: A visual world eye-tracking study. *Cognitive Science* .

DeRose, Keith. 1991. Epistemic possibilities. *The Philosophical Review* 581–605.

Douven, Igor. 2006. Assertion, knowledge, and rational credibility. *The Philosophical Review* 115(4). 449–485.

Edgington, Dorothy. 1995. On conditionals. *Mind* 104(414). 235. doi:10.1093/mind/104.414.235.

Edgington, Dorothy. 1997. Vagueness by degrees. In Rosanna Keefe and Peter Smith (eds.), *Vagueness: A reader*, 294–316. MIT Press.

Egan, Andy. 2007. Epistemic modals, relativism and assertion. *Philosophical Studies* 133(1). 1–22.

von Fintel, Kai. 2001. Counterfactuals in a dynamic context. In Michael Kenstowicz (ed.), *Ken hale: A life in language*, 123–152. MIT Press.

von Fintel, Kai and Anthony Gillies. 2008. CIA leaks. *The Philosophical Review* 117(1). 77.

von Fintel, Kai and Anthony Gillies. 2010. Must... stay... strong! *Natural Language Semantics* 18(4). 351–383.

Franke, Michael and Tikitu de Jager. 2007. The relevance of awareness. In Paul Dekker Maria Aloni and Floris Roelofsen (eds.), *Sixteenth Amsterdam Colloquium*, 97–102. ILLC/Department of Philosophy, University of Amsterdam.

Franke, Michael and Tikitu de Jager. 2011. Now that you mention it: Awareness dynamics in discourse and decisions. In *Language, games, and evolution*, 60–91. Springer.

Giannakidou, Anastasia. 1999. Affective dependencies. *Linguistics and Philosophy* 22(4). 367–421.

Ginzburg, Jonathan. 1995a. Resolving questions, I. *Linguistics and Philosophy* 18(5). 459–527.

Ginzburg, Jonathan. 1995b. Resolving questions, II. *Linguistics and Philosophy* 18(6). 567–609.

Glymour, Clark N. 2001. *The mind's arrows: Bayes nets and graphical causal models in psychology*. MIT press.

Gopnik, Alison and Laura Schultz (eds.). 2007. *Causal learning: Psychology, philosophy, and computation*. Oxford University Press.

Groenendijk, Jeroen and Martin Stokhof. 1984. Studies in the Semantics of Questions and the Pragmatics of Answers. Ph.D. thesis, University of Amsterdam.

Hacking, Ian. 1967. Possibility. *The Philosophical Review* 143–168.

Hawthorne, John. 2004. *Knowledge and lotteries*. Oxford University Press.

Heit, E. and C.M. Rotello. 2010. Relations between inductive reasoning and deductive reasoning. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 36(3). 805.

de Jager, Tikitu. 2009. "Now That You Mention It, I Wonder..."': Awareness, Attention, Assumption. Ph.D. thesis, University of Amsterdam.

Karttunen, Lauri. 1972. Possible and must. In John Kimball (ed.), *Syntax and semantics*, vol. 1, Seminar Press.

Klecha, Peter. 2012. Positive and conditional semantics for gradable modals. In Anna Chernilovskaya Ana Aguilar Guevara and Rick Nouwen (eds.), *Sinn und bedeutung 16*, vol. 2, 363–376. MIT Press.

Klecha, Peter. 2014. *Bridging the Divide: Scalarity and Modality*: University of Chicago dissertation.

Knobe, Joshua and Seth Yalcin. 2014. Epistemic modals and context: Experimental data. *Semantics*

and Pragmatics 7. 1–21.
Koller, Daphne and Nir Friedman. 2009. *Probabilistic graphical models: Principles and techniques*. The MIT Press.
Kratzer, Angelika. 1991. Modality. In Arnim von Stechow and Dieter Wunderlich (eds.), *Semantics: An international handbook of contemporary research*, 639–650. de Gruyter.
van Kuppevelt, Jan. 1995. Discourse structure, topicality and questioning. *Journal of Linguistics* 31(01). 109–147.
Lackey, Jennifer. 2007. Norms of assertion. *Noûs* 41(4). 594–626.
Lassiter, Daniel. 2010. Gradable epistemic modals, probability, and scale structure. In Nan Li and David Lutz (eds.), *Semantics & Linguistic Theory (SALT) 20*, 197–215. CLC Publications.
Lassiter, Daniel. 2011. Measurement and Modality: The Scalar Basis of Modal Semantics. Ph.D. thesis, New York University.
Lassiter, Daniel. 2014. Modality, scale structure, and scalar reasoning. *Pacific Philosophical Quarterly* 95(4). 461–490.
Lassiter, Daniel. 2015. Epistemic comparison, models of uncertainty, and the disjunction puzzle. *Journal of Semantics* 32(4). 649–684.
Lassiter, Daniel. 2016. *Graded Modality: Qualitative and Quantitative Perspectives*. To appear with Oxford University Press.
Lassiter, Daniel and Noah D. Goodman. 2013. Context, scale structure, and statistics in the interpretation of positive-form adjectives. In Todd Snider (ed.), *Semantics & Linguistic Theory (SALT) 23*, 587–610. CLC Publications.
Lassiter, Daniel and Noah D. Goodman. 2015a. Adjectival vagueness in a Bayesian model of interpretation. *Synthese* .
Lassiter, Daniel and Noah D. Goodman. 2015b. How many kinds of reasoning? Inference, probability, and natural language semantics. *Cognition* 136. 123–134.
Lewis, David. 1979. Scorekeeping in a language game. *Journal of Philosophical Logic* 8(1). 339–359.
Lewis, David. 1996. Elusive knowledge. *Australasian Journal of Philosophy* 74(4). 549–567.
Lyons, John. 1977. *Semantics, volumes I and II*. Cambridge University Press.
MacFarlane, John. 2011. Epistemic modals are assessment-sensitive. In A. Egan and B. Weatherson (eds.), *Epistemic modality*, Oxford University Press.
MacFarlane, John. 2014. *Assessment Sensitivity: Relative Truth and its Applications*. Oxford University Press.
Mills, Elizabeth S. 1999. Working with historical evidence: Genealogical principles and standards. *National Genealogical Society Quarterly* 87(3). 165–84.
Moss, Sarah. 2013. Epistemology formalized. *Philosophical Review* 122(1). 1–43.
Moss, Sarah. 2015. On the semantics and pragmatics of epistemic vocabulary. *Semantics and Pragmatics* 8. 1–81.
Nadathur, Prerna and Daniel Lassiter. 2014. Unless: an experimental approach. In *Proceedings of sinn und bedeutung*, vol. 19, .
Nuyts, Jan. 2001. Subjectivity as an evidential dimension in epistemic modal expressions. *Journal of pragmatics* 33(3). 383–400.

Oaksford, Mike and Nick Chater. 2007. *Bayesian rationality: The probabilistic approach to human reasoning*. Oxford University Press.

Over, David E. 2009. New paradigm psychology of reasoning. *Thinking and Reasoning* 15(4). 431–438.

Palmer, Frank Robert. 1979. *Modality and the English modals*. Longman London.

Pearl, Judea. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann.

Pearl, Judea. 2000. *Causality: Models, reasoning and inference*. Cambridge University Press.

Portner, Paul. 2009. *Modality*. Oxford University Press.

Radford, Colin. 1966. Knowledge: By examples. *Analysis* 27(1). 1–11.

Rips, Lance. 2002. Reasoning imperialism. In Renee Elio (ed.), *Common sense, reasoning, and rationality*, 215–235. Oxford University Press.

Rips, L.J. 2001. Two kinds of reasoning. *Psychological Science* 12(2). 129.

Roberts, Craige. 2012. Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics & Pragmatics* 5. 1–69.

van Rooij, Robert. 2003. Questioning to resolve decision problems. *Linguistics and Philosophy* 26(6). 727–763.

Rudin, Deniz. 2016. Deriving a variable-strength *Might*. To appear in *Proceedings of Sinn und Bedeutung 2015*.

Russell, Stuart and Peter Norvig. 2010. *Artificial Intelligence: A Modern Approach*. Prentice Hall.

Sloman, Steven A. 2005. *Causal models: How we think about the world and its alternatives*. Oxford University Press.

Spirtes, Peter, Clark Glymour and Richard Scheines. 1993. *Causation, prediction, and search*. MIT press.

Stanley, Jason. 2005. Fallibilism and concessive knowledge attributions. *Analysis* 65(286). 126–131.

Stanley, Jason. 2008. Knowledge and certainty. *Philosophical Issues* 18(1). 35–57.

Stone, Matthew. 1994. The reference argument of epistemic *must*. In *First international workshop in computational semantics (IWCS 1)*, 181–190.

Swanson, Eric. 2006. Interactions With Context. Ph.D. thesis, MIT.

Swanson, Eric. 2010. On scope relations between quantifiers and epistemic modals. *Journal of Semantics* .

Swanson, Eric. 2015. The application of constraint semantics to the language of subjective uncertainty. *Journal of Philosophical Logic* 1–26.

Teigen, Karl. 1988. When are low-probability events judged to be 'probable'? Effects of outcome-set characteristics on verbal probability estimates. *Acta Psychologica* 68. 157–174.

Tenenbaum, Joshua B., Charles Kemp, Tom L. Griffiths and Noah D. Goodman. 2011. How to grow a mind: Statistics, structure, and abstraction. *Science* 331(6022). 1279–1285.

Turri, John. 2016. *Knowledge and the norm of assertion: An essay in philosophical science*. Open Book Publishers.

Unger, Peter. 1971. A defense of skepticism. *The Philosophical Review* 80(2). 198–219.

Unger, Peter. 1975. *Ignorance*. Oxford University Press.

Weiner, Matthew. 2005. Must we know what we say? *The Philosophical Review* 227–251.

Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford University Press.

Windschitl, Paul D. and Gary L. Wells. 1998. The alternative-outcomes effect. *Journal of Personality and Social Psychology* 75(6). 1411–1423.

Yalcin, Seth. 2005. A puzzle about epistemic modals. *MIT Working Papers in Linguistics* 51. 231–272.

Yalcin, Seth. 2007. Epistemic modals. *Mind* 116(464). 983–1026.

Yalcin, Seth. 2010. Probability operators. *Philosophy Compass* 5(11). 916–937.

Yalcin, Seth. 2011. Nonfactualism about epistemic modality. In A. Egan and B. Weatherson (eds.), *Epistemic modality*, 295–332. Oxford University Press.

Yalcin, Seth. 2012a. Bayesian expressivism. *Proceedings of the Aristotelian Society* 112. 123–160.

Yalcin, Seth. 2012b. Context probabilism. In M. Aloni, V. Kimmelman, F. Roelofsen, G. W. Sassoon, K. Schulz and M. Westera (eds.), *Logic, language and meaning*, vol. 7218 Lecture Notes in Computer Science, 12–21. Springer.

Yalcin, Seth. 2012c. A counterexample to modus tollens. *Journal of Philosophical Logic* 41(6). 1001–1024.