Hedging in discourse

Peter van Elswyk, Northwestern University

Epistemic terms of various syntactic categories can uniformly be used to do the same thing—to hedge. This essay clarifies hedging as a phenomenon and explains how hedging happens by advancing the positional theory. The guiding idea is that, in uttering declaratives, speakers signal what their epistemic position is towards the content put into play by the declarative. The default signal is that the speaker knows. But when an epistemic term hedges, the term overrides the default. The non-default signal sent is that the speaker or someone else occupies the position indicated by the term. To make that idea precise, the positional theory treats hedging as a discourse function. Terms hedge because of how a declarative containing an epistemic term is situated within a discourse.

Forthcoming in *Synthese*

1 Introduction

A speaker does not always come out and say it. When asked a question like *Which plant wilted?*, speakers often refrain from using a bare declarative like (1) to give an answer.

(1) The monstera wilted.

They do not stay silent either. They hedge.[1] They do so by pinning an epistemic term to a declarative. As (2) through (5) illustrate, the term can be a parenthetical verb like *think* placed at the end of a sentence, an adverb like *probably* attached at the beginning, an auxiliary like *might* stuffed in the middle, or an adjective like *possible* paired with a sentential subject.

(2) The monstera wilted, I think

(3) Probably, the monstera wilted.

(4) The monstera might have wilted.

(5) That the monstera wilted is possible.

The broad result is the same whatever the epistemic term. The hedged declarative provides the same answer—that the monstera wilted—without the same intuitive strength as (1). This is the weakness effect.

It is initially tempting to attribute the weakness effect entirely to the compositional semantics of the epistemic terms. After all, the answer to the question is

[1] The hedging discussed in this essay is sometimes called shield hedging (Prince et al., 1982), illocutionary hedging (Brown and Levinson, 1987; Fraser, 2010), or relational hedging (Prokofieva and Hirschberg, 2014). Such hedging changes the strength or confidence with which a content is stated as opposed to changing the content itself. It contrasts with modifiers like *sorta* or *kinda* as in *The monstera sorta wilted*. See Lakoff (1973), Lasershon (1999), and Morzycki (2011) for discussion of the latter.

provided through the content that is embedded under the term, and the terms denote differing degrees of uncertainty that the speaker has towards that embedded content. Nevetheless, we will get a better explanation of the weakeness effect by looking beyond compositional semantics.

There are two reasons why. First, the weakness effect is not caused by every use of an epistemic term. This is an old point.[2] For example, here is Stalnaker (1970, 286-287):

> A sentence of the form 'It may be that *P*' can be interpreted as expressing a modal proposition, that proposition being a function of *P*, or it may be interpreted as making explicit that the negation of *P* is not presupposed in the context. In the latter case, *P* is the only proposition involved. The modal word indicates the force with which it is expressed. A sentence of the form 'I suppose that *P*' may be meant as a report about a supposition of the speaker, or as a rather tentative assertion of *P*.

To illustrate with our present examples, notice that the question *Which plant wilted?* is not about anyone's evidence or attitudes. But sometimes our questions are (Hacquard and Wellwood, 2012). Imagine that we realize on vacation that we forgot to ask the neighbor to water the houseplants. We are then considering our plants individually and asking, given our shared knowledge about their water needs, whether it is epistemically possible that we will return to find that plant to have wilted.

(6) (a) Might the monstera have wilted?
(b) Yes, the monstera might have wilted.
(c) I think the monstera might have wilted.

(6b) contains the same epistemic term as (4) but the answer it provides is not hedged. The speaker comes right out and answer the question as opposed to answering with embedded content. In this respect, (6b) contrasts with (6c). With its extra layer of embedding, (6c) is hedged with *I think* to answer with *The monstera might have wilted*. What discourse (6) illustrates is that the weakness effect is discourse-sensitive (*c.f.* Section 3). Since the effect is discourse-sensitive, it cannot be owed merely by the compositional semantics. The weakness effect depends on how such terms are related to the discourse.

Second, the weakness effect does not require that the speaker's own uncertainty towards a content be indicated. Someone else's uncertainty can be. Consider response (7b).

(7) (a) Which plant wilted?

[2] Once upon a time, the standard view was that epistemic terms were exclusively force modifiers. Toulmin (1958) and Austin (1962) are early examples. Stalnaker (1970) offers a middleground where force modification is one interpretation that terms can receive. But as force modifier views were abandoned, the uses of epistemic terms that motivated them were neglected. In §6, I discuss force modifier views, and contemporary views like Murray (2017) and Moss (2015, 2018) that offer a middleground.

(b) Theresa thinks the monstera wilted.

Opting for (7b) over the the flat-out *The monstera wilted* conveys less confidence in the answer. It weakly answers via embedded content. But unlike the other examples, (7b) does not entail anything about the speaker whatsoever. Since speakers can hedge without indicating outright anything about their own uncertainty, the compositional semantics of epistemic terms is insufficient to explain the weakness effect (*c.f.* Section 4). The effect depends on the function or role that epistemic terms perform within the broader discourse.

In this essay, I advance the POSITIONAL THEORY to explain hedging. The guiding picture is simple. When a speaker uses a declarative, they signal who takes what epistemic position towards the content put into play by the declarative. The default signal is that the speaker knows. But when an epistemic term hedges, the term overrides the default by sending a different signal. That non-default signal indicates that either the speaker or someone else occupies the position indicated by the term. The weakness effect is produced because the position signaled no longer entails that the speaker knows the content in play.

Though the guiding picture is simple, how to develop the picture into an informative theory is not. An immediate choice is deciding where to locate signaling about positions. I develop the positional theory at the level of discourse. Discourses consist of utterances structured with COHERENCE RELATIONS (Hobbs, 1985; Kehler, 2002; Asher and Lascarides, 2003; Lepore and Stone, 2015; Stojnić, 2021). The discourse function of an utterance depends on the relation(s) it bears to the broader discourse. I will argue that hedging is a discourse function that terms can have within certain structures.

Coherence relations are part of semantics in a broad sense. They have conventional meaning and contribute to the meaning of a discourse. In explaining hedging as a discourse function, I therefore treat it as a semantic phenomenon. The hedging interpretation is a meaning an epistemic term can have at the level of discourse. However, as I understand them, coherence relations are not semantic in the narrower sense that they fall within the purview of the compositional semantics determined by a generative grammar. Coherence relations sit above compositional semantics at a different level of conventional meaning. The positional theory will therefore be compatible with most, if not all, compositional theories of the epistemic terms that can function as hedges.

An approach to discourse on which it is structured by coherence relations differs from one in which it is structured by questions. With the latter approach, every sentence must be understood with reference to a question under discussion (Roberts, 1996/2012; Ginzburg, 2012). Sentences either ask questions or answer them. As a result, discourses exhibit coherence in no other way. In contrast, a relations-based approach to discourse appeals to an array of distinct relations to structure discourse. These relations include a question-answer relation—see §5.3 below for discussion—such that the approach can also structure discourse as a question-answer exchange, but it is not limited to invoking just this one

relation. So discourses can exhibit coherence in a variety of ways. Fully litigating the differences between these approaches is outside the scope of this paper. I mention the difference to be transparent about a theoretical choice that I make in my approach but which is not argued for.

The essay divides in two. In the first part, hedging is brought into sharper focus as something speakers do with epistemic terms. In the second, it is explained. But I can be more specific. In §2, I argue that the weakness effect consists in suspending the signal that the speaker knows. In §3, I show how the hedging interpretation is discourse-sensitive for epistemic verbs, adverbs, auxiliaries, and adjectives. In §4, I similarly show how hedging is perspective-insensitive for such terms. In §5, the positional theory is developed. In §6, alternative explanations are considered before concluding in §7.

2 Knowledge suspension

Knowledge sets the threshold for hedging. An epistemic term can hedge only if it indicates a position towards a content that is weaker than knowledge. Taking a cue from Benton and van Elswyk (2020), the easiest way to observe the threshold is to consider attempts at hedging by indicating knowledge. Compare the following answers.

(8) (a) Which plant wilted?
(b) I think that the monstera wilted.

(9) (a) Which plant wilted?
(b) I know that the monstera wilted.

The answer in (8b) is hedged. The speaker says only what they think, which is weaker than if they had replied with the bare *The monstera wilted*. However, the answer in (9b) is not hedged. There is an interpretation of (9b) where it expresses uncertainty about what *other* plants wilted. But there is no uncertainty conveyed about whether the monstera wilted.

To explain why knowledge is the threshold, I adopt this widely endorsed generalization.[3]

> KNOWLEDGE HYPOTHESIS (K-hypothesis)
> For a speaker S and bare declarative d expressing at-issue content p in a context c, S's use of d signals S as knowing p in c.

Let's call the signal posited by the hypothesis the *κ-signal*. I do not take a stance on how to explain the signal because the positional theory is compatible with various explanations. However, it will be instructive at various junctures to consider

[3] See Williamson (2000), Adler (2002), Blaauw (2012), Benton (2011, 2012, 2016a,b, forthcoming), DeRose (2002, 2009), Kelp (2018), Reynolds (2002), Sutton (2005), Schaffer (2008), Simion (2016), Turri (2010, 2011, 2013), and citations therein.

particular explanations. When I reach those points, I will illustrate with both assertoric and semantic explanations.[4]

The explanation for why knowledge is the threshold is straightforward given the κ-hypothesis. Knowledge is the default epistemic position associated with a declarative. Explicitly indicating that one knows cannot override this default. But a term indicating an epistemic position that is weaker than knowledge can override. When it does, the term hedges. Accordingly, the weakness effect consists in suspending the κ-signal by sending a different signal—namely, that the speaker or someone else occupies a weaker position.

The suspension of the κ-signal is noticeable by considering the data that motivates the κ-hypothesis. When a declarative with an epistemic term interpreted as a hedge replaces the bare declarative in such data, the intuitive judgments of the data reverse. Judgments either flip from regarding the data as being infelicitous to regarding it as felicitous or *vice versa* depending on what the judgment was initially. Since the initial judgment of felicity or infelicity motivated the presence of the κ-signal, the flipped judgment motivates the the κ-signal's absence. In this section, I will illustrate such reversals only with parenthetical verbs. Discussion in §3 expands to include other terms.

Consider the peculiar conjunction observed by Moore (1942, 1962). Moorean conjunctions have two ingredients: a declarative d with content p anchored to the perspective of a subject $\mathcal{S}$ and a disavowal that $\mathcal{S}$ knows the primary content of d. Though p and a disavowal of S's knowledge of p are truth-conditionally compatible, Moorean discourses ring as if outright contradictions.

(10) (a) Which plant wilted?
(b) # I don't know. But the monstera wilted.

(11) (a) Which plant wilted?
(b) I don't know. But I think the monstera wilted.

(10b) shows a variant of the discord. The defectiveness of Moorean conjunctions is predicted by the κ-hypothesis. If *The monstera wilted* signals that the speaker knows the monstera wilted, then disavowing such knowledge will contradict the κ-signal. However, (11b) shows that the defectiveness disappears in the presence of a parenthetical. Since the presence of the κ-signal explains why Moorean conjunctions are infelicitous, the signal's absence explains why conjunctions with hedged declaratives are felicitous.

[4] Though the κ-signal is commonly associated with the knowledge norm of assertion, there is surprising flexibility in explaining how a declarative signals the speaker's attitude. On illocutionary explanations where a declarative performs an assertion or a similar speech act, the signal could also be owed to intention-recognition, an extra-semantic convention, a social norm, or some combination thereof. See Murray and Starr (2018) for a helpful survey of different illocutionary mechanisms. On a semantic explanation, the signal about the speaker's attitude might be a owed to the declarative mood (Gutzmann, 2015), a covert parenthetical verb (van Elswyk, 2021), a covert assertion operator (Chierchia, 2006; Hacquard, 2006; Meyer, 2013), or the broader architecture of the semantic theory when it is dynamic (Rett, 2021).

The second kind of evidence for the κ-hypothesis is challenge data noted by Unger (1975). Participants not willing or reluctant to accept what was stated may challenge the speaker. A common challenge is a question like (12b) that presupposes the speaker knows.

(12) (a) The monstera wilted.
(b) How do you know?

(13) (a) The monstera wilted, I think.
(b) # How do you know?

The challenge is natural. The κ-hypothesis explains why. The use of a bare declarative signals that the speaker knows. So a participant can ask for elaboration on why the speaker was entitled to signal that they know. But the naturalness reverses when the declarative contains a hedging term like a parenthetical verb. The signal's absence explains the unnaturalness of these challenges for hedged declaratives. Since hedged declaratives do not send the κ-signal, participants cannot ask how the speaker knows.

Further examples abound because all data motivating the κ-hypothesis reverses. But since my aim is no to defend a well-established hypothesis, discussion of these two lines of data is sufficient to show that the weakness effect consists in suspending the κ-signal, and that the suspension is achieved by replacing the κ-signal with the signal that the speaker or someone else occupies a weaker position towards the content.[5]

3 Discourse-sensitivity

Epistemic terms come from a variety of syntactic categories. To make the scope of this essay manageable, I limit attention to the following four categories of terms in English.

Verb	*I think, I heard, I believe, I guess…*
Adverb	*probably, maybe, perhaps, possibly, reportedly…*
Auxiliary	*might, can, must, may, can't…*
Adjective	*possible, probable, believable, alleged…*

Table 1: Epistemic terms

An epistemic term is typically theorized as relating an underlying proposition to a body of information. The underlying proposition is the PREJACENT. I will argue that whether an epistemic term can hedge by suspending the κ-signal depends on how the prejacent is related to the discourse.

[5] Whether Moorean conjunctions and challenge data support the κ-hypothesis because it well-explains them is disputed. See Lackey (2007), Kvanvig (2009), and McKinnon (2012). Among others, Turri (2010), Benton (2016b), and van Elswyk (2021) offer replies. More recently, Mandelkern and Dorst (2022) have attempted to problematize this very data. But consult van Elswyk and Benton (2023) for a reply.

Assuming that questions are sets of propositions in the style of Hamblin (1973), an answer entails the truth or falsity of at least one proposition in a question's denotation. Question-reply exchanges therefore offer a precise way to distinguish ways in which a declarative can be related back to the previous discourse (Simons, 2007). In what follows, I will consider question-reply exchanges where the whole content of the declarative hosting the term is an answer, and exchanges where the prejacent alone constitutes an answer. I will show that terms hedge only when the term's prejacent is what answers.

To show that the κ-signal is suspended, the question-reply exchanges considered will feature a reply that is the variant the Moorean conjunction provided in §2. That reply begins with *I don't know* and is followed by a declarative with an epistemic term.[6] If the declarative that follows is compatible with the initial disavowal of knowledge, the compatibility indicates the κ-signal is suspended. In contrast, incompatibility with the knowledge disavowal indicates the κ-signal remains present.

3.1 Verbs

Non-factive verbs are the best place to look for verbs that do not indicate knowledge. Candidates include attitude verbs *think* and *guess*, evidential verbs like *seems* and *heard*, and speech act verbs like *wager*, *conjecture*, and *propose*. Such verbs can surface in two positions. They can be the obligatory verb in a declarative's matrix clause, or optionally attach as a parenthetical. Verbs in either position can be interpreted as hedges. Let's start with matrix verbs.

Compared to a bare declarative, declaratives hosting a non-factive verb are intuitively weaker. As (14b) illustrates, they can be preceded by a disavowal of knowledge without discord.

(14) (a) Which plant wilted?
(b) I don't know. But I think that the monstera wilted.

(15) (a) Do you think that the monstera wilted?
(b) # I don't know. But I think that the monstera wilted.

But they do not always hedge. As advertised, the hedging interpretation is unavailable when the verb contributes to an answer for the question under discussion. In addition to not being intuitively weaker like (14b), the Moorean conjunction displayed in (15b) is defective.

Verbs in a parenthetical position are different. Like matrix verbs, they hedge in discourses where the prejacent alone constitutes an answer to a prior question. As

[6] The benefit of considering this variant of Moorean conjunction is that it allows us to diagnose whether the κ-signal is present in a discourse-sensitive manner. It does so by not relying on propositional anaphora. For contrast, consider *The monstera might have wilted, but I don't know that*. Whether this conjunction is felicitous depends on how one resolves the denotation of *that*. A lot of factors may play into this resolution. Discourse structure is one of them, but arguably not the only one. But *I don't know* exploits null complement anaphora to the prior question. So it always depend on discourse.

Hooper (1975, 101) noted awhile ago, "a parenthetical qualifies... by suspending the implication that the speaker knows the proposition to be true." (16b) illustrates the point.

(16) (a) Which plant wilted?
(b) I don't know. But the monstera wilted, I think.

And yet, it is very difficult to eliminate the hedging interpretation. When we try to construct a question-reply exchange where the parenthetical verb contributes to an answer like the matrix verb did in (14b), the resulting exchange is noticeably degraded.

(17) (a) Do you think the monstera wilted?
(b) ? The monstera wilted, I think.

A reply like (17b) does not display a competent use of language. It is easy to recover an intended meaning where the verb is interpreted as if it were in matrix position, but the recovery process is prompted by infelicity. The fact it is infelicitous supports the claim in Urmson (1952, 484) that "the whole point of some parenthetical verbs is... to weaken the claim to truth which would be implied by a simple assertion." While it is easy to find a discourse in which parenthetical verbs hedge, discourses that create ordinary interpretations for other terms are defective with parenthetical verbs.[7]

3.2 Adverbs

Epistemic adverbs provide many options for indicating a position weaker than knowledge. Examples include attitude adverbs like *presumably* and *supposedly*, evidential adverbs like *reportedly* and *seemingly*, and modal adverbs like *probably* and *possibly*. An answer to a question fronted by an epistemic adverb lacks the same confidence as a bare answer. A fronted reply like (18b) is also felicitous in a Moorean conjunction.

(18) (a) Which plant wilted?
(b) I don't know. But probably the monstera wilted.

(19) (a) Which plant probably wilted?
(b) # I don't know. But the monstera probably wilted.

The hedging interpretation vanishes once we change the question in (19a) to concern what is probable. The reply in (19b) is not as weak and the Moorean continuation now yields infelicity. So the hedging interpretation is discourse-sensitive for adverbs.

[7] This difference between matrix and parenthetical verbs with respect to the mandatoriness of the hedging interpretation is addressed directly or indirectly by the semantic literature on parentheticals. See Asher (2000), Jayez and Rossari (2004), Murray (2017), Koev (2019), and van Elswyk (2021), for example. The theory defended here is, in principle, compatible with each of these views because it is neutral on the semantics.

3.3 Auxiliaries

Epistemic auxiliaries display less semantic variety than verbs and adverbs. Examples include *might*, *can*, and *could*. The reply in (20b) demonstrates that they hedge when the prejacent alone is an answer. By hosting an epistemic auxiliary, the declarative uttered is both intuitively weaker and felicitous in a Moorean conjunction.

(20) (a) Which plant wilted?
(b) I don't know. But the monstera might have wilted.

(21) (a) Might the monstera have wilted?
(b) # I don't know. But the monstera might have wilted.

But change the question, and change whether the term hedges in the discourse. In stark contrast to (20b), the reply in (21b) is not weaker nor is the Moorean conjunction felicitous. So the availability of the hedging interpretation for epistemic auxiliaries is also discourse-sensitive.

Missing is mention of *must*. The current literature reports diversity in opinion about whether a declarative with *must* like *The monstera must have wilted* is as strong as or weaker than a bare declarative.[8] The generalization being defended—that the hedging interpretation of epistemic terms is available in question-reply exchanges only when the prejacent is an answer—is, in principle, compatible with *must* being strong or weak. The generalization would just manifest differently. If *must* is weak such that it does not entail its prejacent, it will hedge in question-reply exchanges where only its prejacent is an answer.[9] If *must* is as strong as a bare declarative, it will not be interpreted as such.

I regard *must* as strong for the reasons and defense given by von Fintel and Gillies (2010, 2021). Corroborating evidence comes from our diagnostic for hedging. Exchange (22) demonstrates. Setting aside judgment about weakness because they vary, note that the Moorean continuation produces an infelicitous discourse in (22b).

(22) (a) Which plant wilted?
(b) # I don't know. But the monstera must have wilted.

So *must* does not hedge in the discourse configuration that verbs, adverbs, and auxiliaries indicating a position weaker than knowledge do. What explains why is its strength. However, given the diversity of judgment reported on *must*, mileage may vary about (22b).

[8] For discussion, see Kartunnen (1972), Kratzer (2001), von Fintel and Gillies (2010, 2021), Matthewson (2015), Lassiter (2016), Giannakidou and Mari (2016, 2018), Goodhue (2017), Mandelkern (2019), and Pinal and Waldon (2019).

[9] Depending on how one explains the κ-signal, *must* may be weak but still not be able to hedge. In particular, bare declaratives will be presuppositionally stronger *á la* Schlenker (2012) if the κ-signal is owed to a factive operator in their logical form (Meyer, 2013; van Elswyk, 2021). The possibility that *must* is logically strong but presuppositionally weak is one that has not been explored in the current literature.

3.4 Adjectives

Epistemic adjectives have similar meanings to the other terms. Examples include attitude adjectives like *supposed* and *presumptive*, evidential adjectives such as *alleged* and *apparent*, and modal adjectives like *possible* and *likely*. Where they differ from the other categories is that they are usually outscoped by tense, quantifiers, and the like because of their syntactic position. As a result, isolating a prejacent to be a standalone answer is more difficult. I revisit this difficulty in §4.3. However, a prejacent can be easily isolated if we pair an epistemic adjective in a verb phrase with a subject that denotes a proposition.

Then what I have been calling the prejacent is provided by a dedicated clause or expression. Examples are below with *that*-clauses. In (23), the *that*-clause is in subject position, and in (24) the *that*-clause has been moved to the right edge of the sentence via *it*-extraposition (Quirk et al., 1985).

(23) That the monstera wilted is likely.

(24) It is likely that the monstera wilted.

Once we are working with either syntactic configuration, the discourse-sensitivity observed with other epistemic terms can be observed with the adjectives. Discourses (25) and (26) show that the hedging interpretation jumps out when the prejacent alone is an answer but disappears when the adjective contributes to an answer.

(25) (a) Which plant wilted?
(b) I don't know. But it is likely that the monstera wilted.

(26) (a) Is it likely that the monstera wilted?
(b) # I don't know. But it is likely that the monstera wilted.

I conclude that epistemic adjectives vindicate the discourse-sensitivity of the hedging interpretation in their own way. When the prejacent can be isolated such that it is able to be a standalone answer to the prior question in a discourse, epistemic adjectives hedge.

3.5 Iterated terms

We have seen that epistemic terms uniformly hedge regardless of their syntactic category when just the term's prejacent is an answer to the prior question. But a matter to explore is what happens interpretively when a sentence contains multiple terms with one embedding the other.

The generalization remains intact. Hedging happens when a prejacent alone is an answer. It just depends on which epistemic term's prejacent is the answer. To illustrate, let's consider what I call a HARMONIC HEDGE. When one term embeds another, and it is the prejacent of the embedded term that is the answer, the terms together hedge.

(27) (a) Which plant wilted?

(b) I don't know. But I think the monstera might have.

As a reply, (27b) is not naturally interpreted as a statement about what epistemic possibilities the speaker thinks or believes. That is perhaps its compositional meaning; it is not how we understand what the speaker is doing. Instead, (27b) is interpreted as a hedged reply that the monstera wilted. Such a harmonic interpretation confirms our generalization that epistemic terms hedge when they are not contributing to an answer.

Harmonic hedging was first observed by McCready (2015). But McCready further argues that iterated terms either have the hedging interpretation harmoniously or not at all. The data does not bear this out. An embedded term can contribute to an answer while the outer term can hedge. For an example, assume the same context for (6). We are considering our plants individually and asking, given our shared knowledge, whether it is epistemically possible that we will return to find that it wilted. The reply (28b) is given. Here the speaker is hedging a statement about what is epistemically possible for them as opposed to hedging harmoniously with both terms.

(28) (a) Might the monstera have wilted?
(b) I don't know. But I think it might have.

Iterated terms therefore further confirm the discourse-sensitivity of the hedging interpretation. When situated in a question-rely exchange, a term hedges when a term does not contribute to an answer, but its prejacent or perhaps its prejacent's prejacent does.[10]

3.6 A digression

We have now seen with a litany of examples that discourse structure influences whether an epistemic term is interpreted as a hedge. In question-reply exchanges, terms hedge only when they fail to contribute to the question's answer because its prejacent alone is the answer. Many readers will notice that whether or not the term contributes to an answer also determines whether the term contributes to content that is AT-ISSUE or NOT-AT-ISSUE. Indeed, whether content is an answer in a question-reply exchange is a standard diagnostic for whether that content is at-issue (Tonhauser, 2012). So a natural question to address at his juncture is how much the discourse-sensitivity of the hedging interpretation is owed to the at-issue/not-at-issue status of the relevant content.

[10] These observations do not immediately bear on whether iterated epistemic terms have or lack a harmonic reading when they are not being used to hedge. Since the discursive explanation of the hedging interpretation I offer in §5 is broadly compatible with different perspectives on whether non-hedging uses are harmonic, I do not see this data as immediately relevant. However, the observations do urge a methodological lesson. In considering what meanings iterated terms have, one has to be careful to consider instances of iteration where none of the terms are hedges. Otherwise one may be considering disjoint data. For various perspectives on iterated epistemic terms, see Yalcin (2007), Anand and Brasoveanu (2010), Huitink (2012), Willer (2013), Giannakidou and Mari (2018), and Moss (2015, 2018).

The answer: *very little*. Though an epistemic term can be characterized as not-at-issue when it hedges, the hedging interpretation being discourse-sensitive is not the same as it being not-at-issue. The reason why is that at-issue status is a separate matter altogether. Many terms can be not-at-issue without being interpreted as hedges. For example, sentence-medial nominal appositives like the *A plant native to Southern Mexico* in *The monstera, a plant native to Southern Mexico, wilted* is not-at-issue. But the nominal appositive does not hedge. It contributes additional content without suspending the κ-signal. As such, an explanation of hedging as a discourse-sensitive phenomenon will need to do more than account for why the epistemic terms can be glossed as not-at-issue.

Until now, I have avoided the at-issue/not-at-issue terminology. I now return to avoiding it. Discussing it would not clarify the data as much as add an unnecessary layer of theory over it. At present, I also do not think that the distinction is well-understood.[11] In most settings, being at-issue is a proxy for possessing a cluster of other properties (*e.g.* being available for propositional anaphora, being an answer to the question under discussion, being the attachment site for coherence relations, being projective content). I will talk about these other properties when they are relevant.

4 Perspective-insensitivity

In each of the previous examples, the term was anchored to the perspective of the speaker either explicitly with the first-person pronoun *I* or implicitly. Consequently, the examples indirectly support the further generalization that the hedging interpretation is perspective-sensitive. To hedge, the term must be about the speaker's attitudes or evidence. Assuming perspective-sensitivity, a natural view emerges. A speaker overrides the default signal that they know by indicating the position short of knowledge that they occupy instead.

But this is a bug as opposed to a feature of the examples. The hedging interpretation does not require the term to be anchored to the speaker's perspective. It can be anchored to any perspective. To hedge, what is required of the signaled position is just that it not entail that the speaker knows. This is how the κ-signal is suspended in a discourse. A straightforward way to accomplish this is to signal that the speaker occupies a weak position. But signaling that a third-party occupies a weak position accomplishes this too.[12]

We can see how signaling that a third-party occupies a weaker position suspends κ-signal by looking at new data. This time, I sort terms according to whether

[11] See Tonhauser et al. (2013), Murray (2014), Syrett and Koev (2015), Hunter and Asher (2016), Snider (2017), and Korotkova (2020) for discussion. Of these, the theory in Hunter and Asher (2016) most resonates with the approach to discourse structure adopted in this essay.

[12] Signaling that a third-party occupies a strong position by using a factive verb will not suspend the κ-signal because of the verb's presupposition. Sentences like *Theresa knows the monstera wilted* still signal that the speaker knows that the monstera wilted via *know*'s presupposition. See Garcia-Carpintero (2020), van Elswyk (2021), and van Elswyk and Benton (2023) for relevant discussion.

they require an overt subject or if the perspective dependence happens implicitly. In the first category are verbs. Whether in matrix or parenthetical position, verbs with a third-person subject can hedge. An example with a matrix verb is provided by (29b).

(29) (a) Which plant wilted?
(b) I don't know. But Theresa said that the monstera wilted.

(30) (a) Which plant did Theresa say wilted?
(b) # I don't know. But Theresa said that the monstera wilted.

Discourses (29) and (30) show the discourse-sensitivity catalogued earlier but for *Theresa said* being interpreted as a hedge. It is straightforward that *Theresa said that the monstera wilted* does not entail that the speaker knows the same content. Nevertheless, it still suspends the κ-signal.

In the next category are adverbs, auxiliaries, and adjectives that are not anchored to a perspective via an overt element. To show that these terms can hedge while being anchored to a perspective other than the speaker's, we need to focus on EXOCENTRIC INTERPRETATIONS (Egan et al., 2004; Dorr and Hawthorne, 2012). Epistemic terms are typically interpreted as being anchored to the speaker or to a group that at least includes the speaker. But exocentric interpretations that anchor the term to the attitudes or evidence of individuals other than the speaker are possible. Stephenson (2007) notes that epistemic terms embedded under attitudes are interpreted as reporting the position of the immediate subject. Accordingly, we can force exocentric interpretations by using restrictor phrases involving attitudes anchored to a third-party's perspective.

The examples below show that the terms still hedge even when the exocentric reading is forced by anchoring the term to the neighbor's perspective. Discourses (31) and (32) show that an adverb anchored to the neighbor's perspective hedges in the right discourse.

(31) (a) Which plant wilted?
(b) I don't know. But given what the neighbor$_i$ knows, the monstera probably$_i$ wilted.

(32) (a) Which plant probably wilted given what the neighbor knows?
(b) # I don't know. But given what the neighbor$_i$ knows, the monstera probably$_i$ wilted.

Next, discourses (33) and (34) show that an auxiliary anchored to the neighbor's perspective hedges depending on whether the prejacent alone is an answer to the prior question.

(33) (a) Which plant wilted?
(b) I don't know. But in view of what the neighbor$_i$ saw, the monstera might$_i$ have wilted.

(34) (a) Which plant might have wilted in view of what the neighbor saw?

(b) # I don't know. But in view of what the neighbor$_i$ saw, the monstera might$_i$ have wilted.

Finally, (35) and (36) show that adjectives patterns with adverbs and auxiliaries in hedging in a discourse-sensitive manner even while still receiving an exocentric interpretation.

(35) (a) Which plant wilted?
(b) I don't know. But according to the neighbor$_i$, it is likely$_i$ that the monstera wilted.

(36) (a) Which plant likely wilted according to the neighbor?
(b) # I don't know. But according to the neighbor$_i$, it is likely$_i$ that the monstera wilted.

These examples round out our survey of epistemic terms. Verbs, adverbs, auxiliaries, and adjectives can all be used to hedge even when they are about the epistemic position of someone else.

One feature of hedging by way of signaling how others are positioned is that it requires more disambiguation than hedging by signaling how one is personally positioned. When hedging through self-attribution, there is no unclarity about what the speaker's own position is. By shifting the perspective to a third-party, there is unclarity. For example, the sentence *Theresa said the monstera wilted* contributed to a hedging interpretation in earlier examples. But it can also be followed with *And she is right* or *And I know that to be true*. In such instances, the third-party reporting does not contribute to a hedging interpretation. This is because the extra conjunct clarifies that the speaker's own position is not weaker than knowledge. Typically, though, hedging by signaling how others are positioned exploits the unclarity to suspend the κ-signal. We learn something indirectly about how the speaker is positioned by learning something explicitly about how someone else is positioned toward the relevant proposition. If they are reported as having a position weaker than knowledge, the speaker is taken to occupy a similar position. In the examples above, the hedging interpretation is encouraged with the conjunct *I don't know*. This conjunct is useful as our Moorean diagnostic (§2). However, it is not strictly necessary for the hedging interpretation.

It might initially seem puzzling that speakers can hedge by signaling how someone else is positioned. But this is a natural reflection of our epistemic predicament. Much of our evidence comes from what other people tell us in conversation. When we do not take ourselves to know what we have testimonial evidence for, we can hedge by being transparent about who are source was. This is what hedging with a term anchored to a perspective other than the speaker's enables.[13]

[13] In this respect, hedging with terms anchored to a non-speaker perspective mirrors the effects of using reportative evidentials. Reportatives are considered weak because they are compatible with a subsequent denial of or denial of belief in the reported content (AnderBois, 2004; Murray, 2017). In working out the semantics of reportatives, many have suggested that their weakness is owed to anchoring the source to a third-party as opposed to the speaker as with perceptual evidentials (Faller, 2019; Bhadra, 2020).

5 The positional theory

The positional theory can now be advanced to explain how epistemic terms suspend the κ-signal in manner that is a discourse-sensitive but perspective-insensitive. I first introduce an approach to discourse where it consists in segments connected by coherence relations (§5.1). Then I develop the positional theory of hedging within such an approach to discourse (§5.2). Finally, I show how the positional theory predicts when a term hedges: it fits the data from §3 and §4, and makes three extra predictions (§5.3).

5.1 Coherence relations

When two or more utterances are gathered together they constitute a discourse. Discourses are structured. They hang together with coherence. As Grice (1989, 26) put it, conversations "do not normally consist of a succession of disconnected remarks, and would not be rational if they did." Theories of discourse structure differ according to how they structure or organize discourses to explain their coherence. Just as a theory of grammar explains how a sentence is constructed for it be grammatical, a theory of discourse structure explains how a discourse can be constructed for it to be coherent.

I adopt an approach that takes discourses to be structured by COHERENCE RELATIONS between segments.[14] To illustrate, consider discourse (37). It is natural to interpret (37b) as the explanation for (37a). The cause of disappointment was the wilting.

(37) (a) I was disappointed.
(b) The monstera wilted.

A relational approach takes this interpretation as being driven by a coherence relation that obtains between (37a) and (37b). In particular, the relation EXPLANATION holds between the two segments.

Example (37) illustrate how coherence relations are key to how we interpret discourses. But coherence relations are also central to how we interpret subsentential elements. Examples include verb phrase ellipsis and the meanings of pronouns, tense, quantifiers, and perhaps most context-sensitive vocabulary (Stojnić, 2021). I illustrate with pronouns because the interpretive influence on pronouns is the most rigorously studied. An oft-cited example owed to Smyth (1994) is (38).

(38) (a) Phil tickled Stanley.

[14] A relational approach to discourse structure is commonplace in computer science and adjacent corners of linguistics. For example, see Hobbs (1985), Grosz and Sidner (1986), Mann and Thompson (1988), Carlson and Marcu (2001), Kehler (2002), Webber et al. (2003), Asher and Lascarides (2003), and Wolf and Gibson (2006). A relational approach is less familiar within philosophy. Notable exceptions include Starr (2014), Pagin (2014), Lepore and Stone (2015), Cumming et al. (2017), Cohen and Kehler (2018, 2021), and Stojnić (2017, 2021) where coherence relations are given leading roles in semantic and/or pragmatic explanations.

(b) Liz poked him.

The second segment of the discourse contains the pronoun *him* and the prior segment contains two candidate antecedents. The pronoun is therefore ambiguous. But the ambiguity resolves depending on which coherence relation is understood to link (38a) and (38b). Suppose they are linked by RESULT, the event described in (38b) is an outcome or consequence of the event described in (38a). Then the pronoun resolves to Phil. Liz's poking is understood as retaliation against Phil for what he did to Stanley. In contrast, suppose they are linked by PARALLEL. What is described by (38a) parallels what is described by (38b). Phil does something to Stanley, and what Liz does parallels that. Then the pronoun resolves to Stanley. Similar interpretive influence of coherence relations on pronouns has been confirmed in a variety of studies.[15]

Theories within the relational approach differ in the number of coherence relations posited, and the rules that govern when segments can be related to build out discourses. I take few sides on these issues to ensure that the positional theory is broadly compatible. However, I do adopt the ATTRIBUTION relation (Wolf and Gibson, 2006; Hunter, 2016).[16] This relation asymmetrically attaches α to β because α indicates who takes what epistemic position towards the content of β. The most common use of ATTRIBUTION is representing the structure projected by a third-party report like (39).

(39) Theresa said the monstera wilted.

We can represent (39) as consisting of three segments at the level of discourse: a segment for the primary discourse unit (π), which, in (39), corresponds to the sentence, a segment for the matrix clause *Theresa said*, and a segment for the embedded clause *The monstera wilted*. The ATTRIBUTION relation is what links the matrix to the embedded clause.

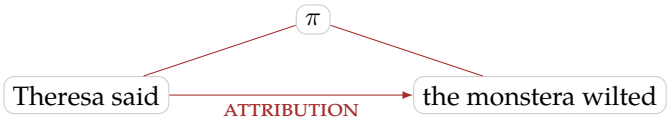

Figure 1: (39)'s structure

Though the common use for ATTRIBUTION is third-party reports, the relation is not limited to third-party reports (Potter, 2019). It can be used to represent self-attribution of one's evidence or attitudes as well. In what follows, I put it to such uses.

[15] For example, see Hobbs (1979), Kehler (2002), Kehler et al. (2008), Kaiser (2009), Mak and Sanders (2013), Kehler and Rohde (2013, 2019), and Hoek et al. (2021). These authors differ on how coherence relations exert interpretive influence, but their experimental results converge on the conclusion that they do with striking consistency.

[16] The semantics of ATTRIBUTION will not matter to my proposal. This will become clear in §5.3. Accordingly, I will not belabor the reader with such details. See also *fn.* 17 below on similar coherence relations.

5.2 Position signaling in discourse

Signaling about epistemic position happens at the level of discourse. Or, so I propose. To make this proposal precise, two commitments are needed. The first is that the κ-signal is represented in discourse. When a speaker utters a bare declarative, the bare declarative projects a structure that represents the speaker as knowing the declarative's primary content.

So far, I have remained non-committal about how to explain the κ-signal. I can continue to be. The suggestion that the κ-signal is represented in discourse does not take a side on what causes the signal *beneath* discourse. To see the neutrality, consider explanations of the κ-signal. On an illocutionary explanation, the κ-signal is a byproduct of a declarative tokening the speech act of assertion. The signal is sent because assertions express or require for their sincerity that speakers know. Such an illocutionary effect gets represented in the structure of the discourse, as illustrated by Figure 2.

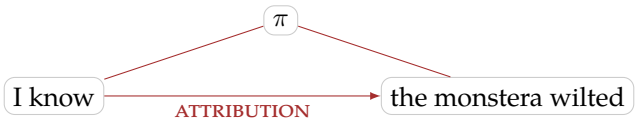

Figure 2: bare structure

With a semantic explanation of the κ-signal, the signal is owed to a covert element in the logical form of a declarative like an assertion operator or a *know*-parenthetical (Meyer, 2013; van Elswyk, 2021). Matters are no different at the level of discourse. That covert element gets represented.[17]

The choice to represent the κ-signal in discourse is justified because the signal contributes to the coherence of a discourse. In particular, it creates attachment sites with which to link the utterance of a declarative to preceding discourse, or sites for subsequent discourse to hook onto as conversation advances. To illustrate, recall from §2 knowledge-presuming challenges like *How do you know?*. The challenges attach to the κ-signal segment, or at least to the bigger segment of which the signal is a part.

[17] Depending on which explanation of the κ-signal one adopts, the proposal that the signal is represented in discourse may require a departure from how the boundaries of discourse structure are sometimes theorized. If it is owed to a covert element, no departure is required. That semantic element will surface in the structure. But if the signal is the effect of an illocutionary act, then a departure is required. For example, Asher and Lascarides (2003) develop segmented discourse representation theory (SDRT) with three levels: informational content, discourse structure, and cognitive modeling. For them, the sincerity conditions of speech acts are represented in cognitive models as opposed to discourse structure. But the κ-signal can be understood as the assumption that an assertion is sincere. Though representing the signal in discourse is a departure, it is still within the spirit of their expansive theory. SDRT deliberately lets coherence relations play many of the roles traditionally played by speech acts.

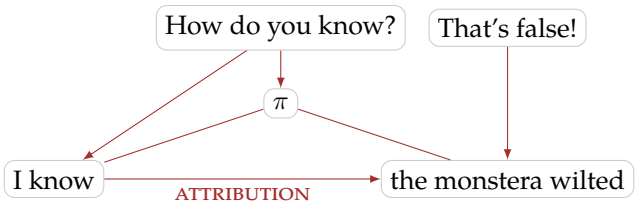

Figure 3: Attachment sites

It elicits an explanation for how the speaker knows the monstera wilted as opposed to an explanation of the wilting itself. For contrast, consider the denial *That's false!*. It only denies that the monstera wilted. As Figure 3 illustrates, the challenge and the denial are best modeled as attaching to different segments of the discourse structure evoked by the bare declarative. Other examples are easily encountered. For most coherence relations, I wager that one can find a discourse where that relation attaches to the segment for the κ-signal, or at least the broader segment to which the κ-signal contributes.

The second commitment required is that epistemic terms typically project a structure with the epistemic term attached to its prejacent with ATTRIBUTION. Some might balk at giving adverbs, auxiliaries, and adjectives their own segment because they are not clauses. But balking is more parochial than substantive. The second commitment is justified for the same reason as the first. The segmentation is needed to explain the coherence of discourse where relations attach to different segments. In particular, the availability of prejacents as attachment sites for denials like *That's false!* clearly justifies this commitment.

In the usual instance of ATTRIBUTION, the primary discourse unit or π corresponds to the whole sentence or to at least a clausal boundary. This was true in Figure 1 for (39). However, taking on board these two commitments requires us to expand our notion of what π is. It cannot correspond with clausal boundaries. Instead, it is unit corresponding to a chunk of discourse that includes the proposition contributed by a clause and an additional satellite connected with ATTRIBUTION that may or not be overtly represented in the sentence uttered. This expanded conception is represented in Figures 2 and 3.

The DISCOURSE FUNCTION of a segment s is determined by the coherence relation that connects s to a structure (Mann and Thompson, 1988; Asher and Lascarides, 2003). For example, when *The monstera wilted* is linked back to *I was disappointed* with EXPLANATION in (37), the function of *The monstera wilted* is to explain. Given the two commitments, the discourse function of the κ-signal and an epistemic term is the same: to attribute, or, more precisely, to identify who stands behind the content of a segment with what epistemic position.

We are now able to precisely characterize position signaling. It is a discourse function. The position signaled for a content is determined by whatever segment in the discourse performs the signaling function for that content. So a segment s carries a signal about the speaker's epistemic position when there is a segment

s' attached to s with ATTRIBUTION or a similar coherence relation.[18] Without such a segment s', no position is signaled. Understood as a discourse function, position signals are not limited to main clause content. Embedded content also carry signals. Recall (39).

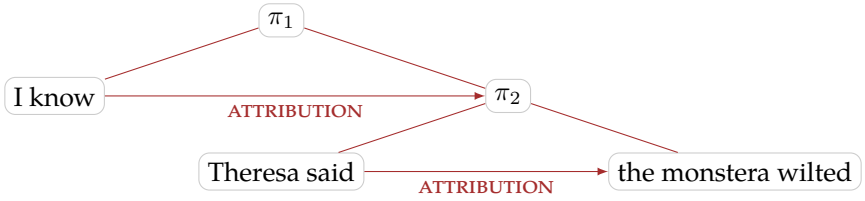

Figure 4: (39) revisited

It was represented with ATTRIBUTION. So the position signaled for the embedded clause was that Theresa gave testimony. But given the commitment that the K-signal surfaces in discourse, the structure of (39) is actually richer. Figure 4 illustrates. The main clause (π_2) receives a position signal from an attached segment just like the embedded clause. A few places in this paper I have described content as being *put into play* by the use of a declarative. The reason why is that sometimes conversational participants are interested in an embedded clause. When they are, the position signaled for that clause can change if it has its own satellite performing this function.

5.3 The weakness effect

Let's put the pieces together. In interpreting what is uttered, participants organize utterances into a discourse with coherence relations. Each new utterance projects its own structure that needs to be attached to the structure of the preceding discourse. But participants face a choice when they encounter a declarative hosting an epistemic term. Since declaratives carry the K-signal and epistemic terms contribute further structure, participants encounter the more elaborate structure of Figure 5.

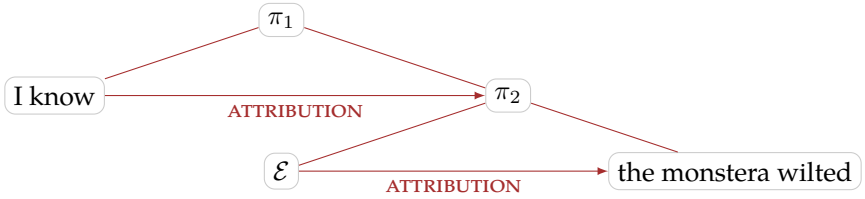

Figure 5: Declarative with epistemic term

[18] I add the proviso *or a similar relation* to make room for other relations that perform the same function of indicating who stands behind a content in what epistemic manner. For example, Hunter (2016) deploys a SOURCE relation and Carlson and Marcu (2001) deploy an EVIDENCE relation in addition to ATTRIBUTION. Since a common complaint with relational approaches to discourse is that they are unconstrained in what relations are posited, investigating how far we can get with ATTRIBUTION is worthwhile.

This structure has three primary attachment sites: the segment to which the κ-signal contributes (π_1), the segment to which $\mathcal{E}$ contributes but the κ-signal does not (π_2), and the prejacent of $\mathcal{E}$. The choice participants face is what attachment site to relate back to the prior discourse.

Hedging happens when the prejacent is the preferred attachment location. Then the epistemic term performs the signaling function as opposed to the κ-signal because the term is the most immediate segment attached with ATTRIBUTION. Put precisely:

> POSITIONAL THEORY
> For any declarative d hosting an epistemic term $\mathcal{E}$, $\mathcal{E}$ hedges a content p in discourse $\mathcal{D}$ if and only if:
>
> (i) $\mathcal{E}$ indicates a position weaker than knowledge in $\mathcal{D}$,
>
> (ii) $\mathcal{E}$ performs the signaling function for the segment with p as its content in $\mathcal{D}$, and
>
> (iii) The segment with p as its content is the preferred attachment site for linking d to the prior structure in $\mathcal{D}$.

What happens to the κ-signal when the prejacent is the attachment site depends on how it gets explained. Suppose an assertoric explanation. Not all uses of declaratives perform assertions. Declaratives with parenthetical verbs are standardly assumed to perform other acts (Williamson, 2000; Garcia-Carpintero, 2004). Perhaps the same goes for all epistemic terms that hedge. Then it is natural to say that the κ-signal just disappears. As a result, there would not be the extra structure dominated by π_1 in Figure 5. On a semantic explanation where the κ-signal is caused by a covert element like an operator or *know*-parenthetical, the added structure will remain. It will just not contribute to the coherence of the discourse by being related to the prior discourse structure with relations of its own. In both cases, the κ-signal is suspended because it no longer plays the signaling function for the content being attended to.

We are now able to explain the data canvassed in §3 and §4. Let's start with discourse-sensitivity. The prior discourse prompts participants to have expectations about how new utterances will cohere. When participants anticipate using relation $\mathcal{R}$, they will prefer the attachment site that works with $\mathcal{R}$. The discourses considered earlier were question-reply exchanges. The relation a participant would expect to use after a question is QUESTION-ANSWER PAIR or QAP. Still assuming a semantics for questions on which they are sets of contents (Hamblin, 1973), QAP obtains between two segments α and β where β is a question when the content of α entails the truth or falsity of at least one content of β. This is how we earlier defined an answer. Participants seeking an attachment site for QAP will therefore be looking for answers. Different questions create different preferences. To see as much, revisit (18) and (19).

(18) (a) Which plant wilted?

(b) I don't know. But probably the monstera wilted.

(19) (a) Which plant probably wilted?

(b) # I don't know. But the monstera probably wilted.

The question in (18a) prompts a preference for QAP to attach to a site whose content entails that a particular plant wilted. In contrast, the question in (19a) creates a preference for QAP to attach to a site whose content entails which plant probably wilted.

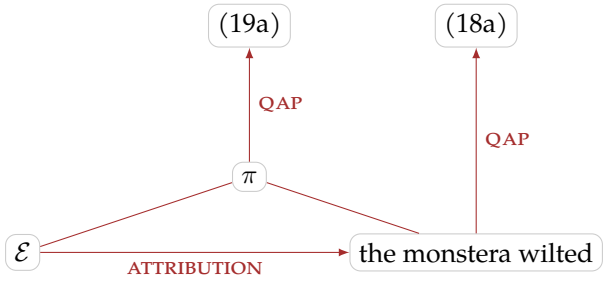

Figure 6: Attachment sites for QAP

Simplifying by excising the κ-signal, Figure 6 illustrates. As predicted by the positional theory, the term *probably* hedges only in (18b). Unlike (19b), (18b) coheres to the prior discourse with QAP attaching to the prejacent because only the prejacent is an answer in the strict sense.

The explanation given for the discourse-sensitivity of adverbs being interpreted as hedges generalizes to the other syntactic categories. What merits extra commentary is iterated terms. With iteration comes more sites for attachment. Again, removing the κ-signal, Figure 7 displays the sites for one term embedding another term.

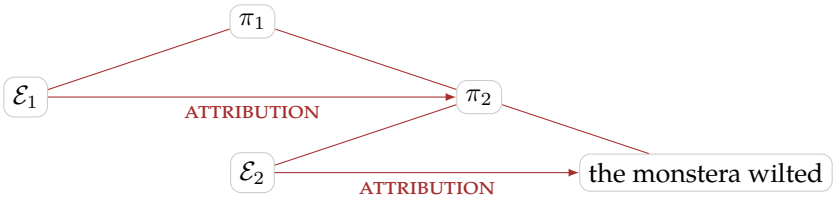

Figure 7: Declarative with epistemic term

A relation can attach to a segment for the whole declarative (π_1), the prejacent for the outer term $\mathcal{E}_1$ (π_2), or the prejacent for the inner term $\mathcal{E}_2$. The predictions made by the attachment-based explanation is that neither of the terms hedge if π_1 is the attachment site, the outer term hedges an epistemic content if π_2 is, and the terms hedge harmoniously if the prejacent is. This battery of predictions is born out by the data in §3.5.

(40) But I think the monstera might have.

For illustration, consider (40) against the backdrop of different questions. The question *Do you think the monstera might have wilted?* will direct attachment to π_1 for non-hedging interpretations of the epistemic terms, *Might the monstera have wilted?* will guide attachment to π_2 to cause the hedging interpretation for the outer term, and *Which plant wilted?* will induce harmonic hedging by privileging the prejacent of the inner term as the attachment site.

Up next is making sense of perspective-insensitivity. The positional theory predicts that a term hedges even if the term is anchored to someone else's perspective. As long as the term is non-factive and performs the signaling function in the discourse structure, the term is eligible for hedging. To illustrate, I repeat example (29) below.

(29) (a) Which plant wilted?
(b) I don't know. But Theresa said the monstera wilted.

The *Theresa said* in (29b) is interpreted as a hedge. The positional theory accounts for why. As Figure 4 depicts, *Theresa said* performs the signaling function for the prejacent and the prejacent is the preferred landing site for attaching backward with QAP. In this way, *Theresa said* hedges because it supplants the K-signal in certain discourse configurations. The same goes for other epistemic terms anchored to a perspective other than the speaker's.

Beyond explaining the data from §3 and §4, the positional theory makes additional predictions about when a term can hedge. I discuss three further predictions. The first is what gets predicted for discourses that are not question-reply exchanges. Since the positional theory is about where a relation attaches and not what relation is installed, hedging should happen in discourses that do not involve QAP. And it does. Suppose you are talking to a friend about Kim, a public figure who knows neither of you, and why she recently appeared disappointed on her social media account. Knowing Kim prizes her rare variegated monstera and that Kim's apartment lost heat in the dead of an Alaskan winter, you utter one of the following discourses.

(41) Kim was disappointed. I think the monstera wilted.

(42) Kim was disappointed. The monstera possibly wilted.

(43) Kim was disappointed. The monstera might have wilted.

(44) Kim was disappointed. It is possible the monstera wilted.

Given that the second sentence uttered contains an epistemic term, the structure evoked by the second sentence contains multiple attachment sites. So the friend faces a choice about how to link the structure evoked by the second sentence to the structure for the preceding sentence. The friend expects to deploy EXPLANATION. So the preferred site is a term's prejacent. Were the friend interpreting your discourse to attach to the segment to which a term contributes, they would arrive at a false and bizarre explanation. Namely, that the reason for Kim's disappointment is that it is

epistemically possible for an individual of which she us unaware that her monstera wilted. In attaching to the prejacent, the positional theory predicts that the term hedges because the term performs the signaling function. Examples (41) through (44) confirm this prediction.

The second prediction concerns declaratives with epistemic terms that project a discourse structure different than the one we have been considering. In §5.2, I took on the commitment that declaratives with epistemic terms typically project a structure where the term is related to the prejacent via ATTRIBUTION. But what is typically the case is not what is always the case. Consider adjectives. As noted in §3.4, isolating a standalone prejacent is often difficult for adjectives. I cheated earlier by focusing on adjectives in verb phrases that occurred with sentential subjects. Then the sentential subject could play the prejacent role. But adjectives can be buried deeper in a sentence. An example is *alleged* in a definite description like *The alleged monstera*. Here the adjective contributes to what plant is denoted and does not hedge.

(45) (a) Which plant wilted?
(b) # I don't know. But the alleged monstera wilted.

(45b) illustrates. I am non-committal regarding how to represent the discourse structure of *The alleged monstera wilted* beyond that its structure does not give a term and a prejacent separate segments. The earlier justification for giving each a distinct segment was that the prejacent could be a distinct attachment site for denials like *That's false!*. But no similar justification can be given for *The alleged monstera wilted*. Denials obligatorily target the content of the whole declarative. Given that *alleged* does not receive its own segment attached with ATTRIBUTION, it cannot perform the signaling function. So the positional theory predicts that it cannot hedge. The term straightforwardly fails condition (ii).

The third prediction concerns what happens in more complicated discourses. So far, only simple examples have been considered where a single relation links a declarative back to the preceding discourse. But more intricate discourses where there are multiple points of attachment are commonplace. When there is a relation attaching a term's prejacent back to the preceding discourse and another attaching back the segment to which the term contributes, the positional theory predicts that a term simultaneously hedges and does not hedge. Such a prediction might seem like a misfire. On the contrary, the prediction is an important vindication. An example is provided by the adjective in (46c). Its prejacent is related to (46a) via QAP. The positional theory correctly predicts that the *possible* hedges; it perform the signaling function for the content that the monstera wilted. However, (46c) is also coherently related to (46b).

(46) (a) Which plant wilted?
(b) A plant in the living room must have wilted.
(c) Therefore it is possible that the monstera wilted.

A speaker is naturally understood as inferring (46c) from (46b), especially given the inference marker *therefore*. These declaratives are plausibly linked by a CONSEQUENCE relation holding between the segments representing the entirety of each declarative.

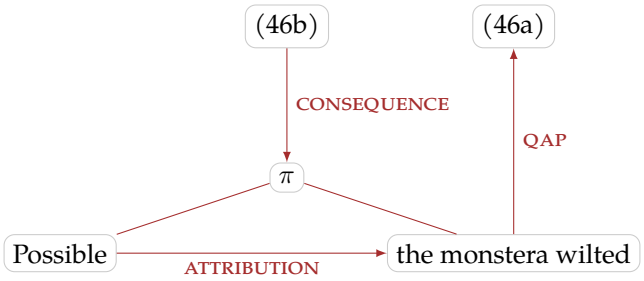

Figure 8: Attachment sites for (46c)

Considered in relation to (46b), *possible* in (46c) does not hedge. Accordingly, the adjective hedges relative to one backwards attachment but not relative to another attachment. The hedging interpretation is not insignificantly discourse-sensitive; it is thoroughly so.

6 Alternative explanations

The positional theory has no developed rivals. Hedging is widely recognized as something speakers do with epistemic terms. But no explanation has been previously offered that generally applies to terms of various syntactic categories, that appreciates the data that terms hedge by suspending the κ-signal (§2), that recognizes that such suspension is discourse-sensitive (§3), and that accommodates that the signal's suspension is perspective-insensitive (§4). Even still, the positional theory and its commitments can be clarified by considering some of the incomplete explanations or claims that have been advanced about hedging and inspecting how they compare.

6.1 Force modification

Insofar as there is a standard view, the view is that hedging is an illocutionary phenomenon (Brown and Levinson, 1987; Bach, 2008; Benton and van Elswyk, 2020). For example, Fraser (2010, 22) writes "There is general agreement today that hedging is a rhetorical strategy, by which a speaker, using a linguistic device, can signal a lack of . . . the full commitment to the force of the speech act being conveyed." Such a view is difficult to assess without a developed account of how a term's presence impacts force. A historical option is that the terms are FORCE MODIFIERS (Dummett, 1973). They take wide-scope over the prejacent to somehow ensure that the speech act performed is illocutionarily weaker than assertion. On the assumption that assertions send the κ-signal, force modification suspends the κ-signal. The problems with a force modifier semantics are numerous. The most serious is that the terms engage in scope configurations that force modifiers cannot.

For example, auxiliaries, adverbs, and adjectives embed under tense, quantifiers, attitude verbs, and questions.[19]

Some might want to avoid this embedding problem by maintaining that epistemic terms are ambiguous between a standard meaning and a force modifying meaning. But it is difficult to see the initial motivation for positing such widespread ambiguity. To explain the data, every epistemic verb, adverb, auxiliary, and adjective that can be used to suspend the κ-signal needs to have at least two meanings. The force modifying meaning will also need to be detectable exclusively in the discourse configurations predicated by the positional theory. Adopting the positional theory is far preferable.

Importantly, the positional theory can be understood as providing a discursive account of force modification. Suppose that the use of a declarative performs a CONSTATIVE SPEECH ACT (Bach and Harnish, 1979; Searle and Vanderveken, 1985). Suppose further that the force of a constative act is determined by what epistemic position it is associated with: assertions are associated with knowledge, conjectures are associated with guesses, and so forth. Then the positional theory will predict which speech acts are performed according to what signal accompanies the content that is the preferred site for backwards attachment. Assertion will be the speech act performed by a bare declarative, but other speech acts are performed when weaker epistemic positions are signaled.

The positional theory does not need to be understood in this manner. One could head in the opposite direction to regard it as dispensing with illocutionary accounts of what speakers do with declaratives in favor of a discursive account. I prefer this take (van Elswyk, 2021). The important point is that a discursive and an illocutionary account are not mutually exclusive. Instead of trying to make force modification happen at the level of compositional semantics, modification can be moved up to the discourse.

6.2 Gricean explanations

Another alternative is that hedging is owed to Gricean calculation. When participants receive a declarative with an epistemic term when they were expecting a bare declarative, the unexpected declarative does something uncooperative. Participants then reason from that uncooperativity to the conclusion that the speaker is using the epistemic term to hedge. McCready (2015) develops such a Gricean explanation for error acknowledgments.

(47) (a) Which plant wilted?
(b) I might be wrong about this, but the monstera wilted.

[19] See Papafragou (2006), von Fintel and Gillies (2007), Swanson (2011), and Hacquard and Wellwood (2012). Even if a force modifier semantics is more plausible for terms like parenthetical verbs than others, the failure of a modifier semantics for adverbs, auxiliaries, and adjectives sinks the received view as a general view. We need an explanation that works for every epsitemic term that hedges, as opposed to just a select few terms.

An example is (47b). Participants confronted with (47b) encounter an utterance that is defective. Since a bare declarative signals that the speaker believes its primary content, (47b) conveys that the speaker simultaneously believes that the monstera wilted and that it possibly did not, or what Yalcin (2007) calls an EPISTEMIC CONTRADICTION. Being charitable in their interpretation, participants seek out the best explanation of what the speaker is doing with their utterance of (47b). The result of this reasoning is that the error acknowledgment is interpreted as an utterance modifier. The acknowledgment makes an utterance of *The monstera wilted* one that is hedged as opposed to asserted.

In this essay, I have not considered error acknowledgments. I do not find (47b) felicitous nor have I found any informants who do. What is preferred to (47b) is an acknowledgment accompanied by a signal about the speaker's epistemic position. The reply in (48b) illustrates. It begins the same, but indicates that the speaker thinks.

(48) (a) Which plant wilted?
(b) I might be wrong about this, but I think the monstera wilted.

However, a reply like (48b) is not defective. It does not saddle the speaker with believing an epistemic contradiction. It conveys that the speaker believes that they think that the monstera wilted while acknowledging it might not have. Though some may find (47b) felicitous, the differences between (47b) and (48b) shows that McCready's explanation of error acknowledgments does not generalize. Hedging is not calculated from epistemic contradictions.

Other Gricean explanations may be attempted that do not prompt participants to calculate an intended meaning from an epistemic contradiction but from the violation of a maxim.[20] Such an explanation is a natural one to pursue if hedging is uncooperative. But it is not. We should resist Gricean overreach, to use the phrase of Cohen and Kehler (2018). It is instructive to compare (49b) and (50b) as replies to the same question.

(49) (a) Which plant wilted?
(b) I don't know.
(50) (a) Which plant wilted?
(b) I don't know. But the monstera might have wilted.

Both convey that the speaker cannot fulfill the expectation of providing an answer that is known. However, (49b) stops there. In contrast, (50b) goes further by

[20] The Maxim of Quantity—the admonition to make contributions as informative as required by the conversation—is somtimes glossed as *assert the stronger* (Jackson, 1979; DeRose, 2002). It might seem that hedging is a way of flagrantly violating this maxim. This is a natural albeit mistaken thought. The notion of strength in this gloss of the Maxim of Quantity concerns the logical strength of the proposition that is contributed to the conversation. For example, it instructs not to say *The monstera or the snake plant wilted* when *The monstera wilted* can be said. The maxiim does not concern how forcefully this proposition is contributed by the speaker. But this latter conception of strength is what is relevant to hedging.

offering information relevant to the question. The contrast illustrates a general feature of hedging. Hedging enables speakers to cooperatively share information without knowing it. They do not have to stay silent nor do they have to mislead participants by signaling that they know when they do not. By hedging, speakers can share relevant information while alerting participants to the epistemic risk associated with that information.

Still more post-Gricean explanations may be suggested that do not require epistemic contradictions nor maxim violations but still appeal to intention recognition to explain hedging. These are natural explanations to reach for when one regards intention-recognition as the general mechanism by which all human communication is accomplished (Bach and Harnish, 1979; Loar, 1981; Schiffer, 1982; Harris, 2019). But nothing about the positional theory is incompatible with such post-Gricean explanations. A common way that hearers recognize speakers intentions is through conventions to which both participants are party. This is an important point about the intention/convention interface advanced early by Strawson (1964). As such, the positional theory can be understood as identifying the conventions with which hearers recognize the speaker's intention to hedge with epistemic terms.

6.3 Semantic explanations

The positional theory does not take a stance on the compositional semantics of epistemic terms that can hedge. As a result, it does not require one to deny this or that account of various terms. For example, it is compatible with expressivist, relativist, and contextualist theories and neutral between static and dynamic theories. I see this as a considerable advantage. But it is natural to wonder whether existing semantic theories can be extended to explain the hedging interpretation. I am skeptical that they can be. To motivate that skepticism, I will offer two cases studies of relevant semantic theories.

The first theory is the update semantics for parenthetical verbs proposed by Murray (2017). On her view, parenthetical verbs make two contributions. First, they update the information state with a secondary, not-at-issue content indicating that the individual specified holds the attitude specified towards the at-issue content. Second, they can weaken how the at-issue is proposed as an update. Murray advertises two modifications. One is to weaken p to $\Diamond p$. The other is to weaken p to no proposed update at all. With the latter modification, the declarative only updates via the secondary proposition.

Though this semantics is limited to parenthetical verbs, parenthetical verbs are exclusively interpreted as hedges (§3.1). An account of parenthetical verb could therefore be treated as a template for how to analyze the hedging interpretation of other terms. At least at a high altitude, extending the theory to other terms yields a plausible picture. When interpreted as hedges, terms would contribute an update with a secondary proposition and can modify how the at-issue is proposed as an update. But implementing this picture will face obstacles. First, a way to shift

terms into a parenthetical-like meaning would need to be devised that applies to terms even when they are anchored to a non-speaker perspective. Second, the shift would need to be induced by discourse structure alone, or the theory would over predict the availability of the hedging interpretation. Third, the data that generally surrounds the semantics of epistemic verbs, adverbs, auxiliaries, and adjectives would need to be accommodated.

Even if the extended version of the theory could overcome these obstacles, it would still be insufficient. Murray's theory only allows parenthetical verbs to weaken the proposed update of the at-issue content in two ways: to $\Diamond p$ or to no proposal at all. But the full range of meanings represented by terms in each of the categories outstrips these two options. For example, probabilistic terms like *probably* or *likely* fall into neither. Accordingly, the theory lacks the resources to capture the full variety of strengths that are exploited.

The next theory is the probabilistic semantics of Moss (2015, 2018). Moss's theory is worth considering because it applies to a wide range of epistemic terms that includes probabilistic ones, and because it distinguishes between two distinct meanings a term can have. A term either receives a thoroughly probabilistic meaning, or a meaning about a contextually determined body of evidence. Moss does not present her account of probabilistic meaning as an account of the hedging interpretation. But perhaps the distinction between thoroughly probabilistic and body of evidence meanings is coextensive with the distinction between hedging and non-hedging interpretations. If it were, Moss's theory would be well-positioned to explain the hedging interpretation if the thoroughly probabilistic meaning could be shown to be discourse-sensitive.

But the distinctions are not coextensive. Instead, they cross-cut. First, the meaning that concerns a contextually relevant body of evidence can still be used to hedge. For Moss, exocentric interpretations are about a contextually relevant body of evidence, but we saw in §4 that terms interpreted exocentrically can still hedge. Second, terms with thoroughly probabilistic meanings can fail to hedge. Moss argues that the thoroughly probabilistic meaning is diagnosed by eavesdropper challengeability. If an eavesdropper can reject the statement merely by having different probabilistic beliefs, the term in the statement is interpreted as having thoroughly probabilistic meaning. Adjusting an eavesdropper case from Egan (2007), we can uncover an example where a term is not interpreted as a hedge while testing positive for being thoroughly probabilistic via the eavesdropper diagnostic. Start with (51b). Its infelicity with the Moorean continuation verifies it is a non-hedging use.

(51) (a) Is it unlikely that James Bond is in London?
(b) # I don't know. But it is unlikely that James Bond is in London.

Next, consider (52b). It is just like (51b) except the Moorean continuation has been removed to reduce noise from the example. It importantly enables an eavesdropper challenge like (52c).

(52) (a) Is it unlikely that James Bond is in London?
(b) Yes, it is unlikely that James Bond is in London.
(c) EAVESDROPPER: No it's not—Bond is almost certainly in London.

What these two considerations show is that when terms have thoroughly probabilistic meanings is not coextensive with nor a proper subset of when terms receive the hedging interpretations. As a result, Moss's semantics for terms receiving a thoroughly probabilistic meaning does not double as a semantics for the hedging interpretation.

7 Conclusion

This essay had two parts. It began by bringing hedging into sharper focus as a function that epistemic terms can perform. §2-§4 contributed to this task by arguing that the weakness effect caused by hedging is owed to suspending the κ-signal, that the suspension of the signal is discourse-sensitive, and that the suspension is also perspective-insensitive. This first part stands alone in individuating hedging as a phenomenon requiring explanation.

The second part in §5-§6 was developing and defending the positional theory. Other alternatives may be envisioned, and I hope they are. But the virtues of the positional theory are straightforward. It explains why epistemic terms of various syntactic categories can be uniformly used to do the same thing—to hedge. It does so without requiring substantial commitments about the compositional semantics of particular terms, or without requiring even a stance on how the use of a bare declarative signals that the speaker knows its primary content. One is free to adopt a traditional explanation on which a bare declarative signals as much by being an act of assertion, or a semantic explanation. By accounting for hedging as a discourse function, the theory treats hedging as a linguistic phenomenon that sits a level above where most theorizing in semantics and pragmatics takes place. In this way, it showcases an alternative way to theorize.[21]

[21] In addition to referees recruited by this journal and others, I am thankful to audiences at Northern Illinois University and University of Chicago, and students in my Fall 2019 and Spring 2021 graduate seminars at University of Wisconsin–Milwaukee. This paper was also supported in Summer 2021 by the UWM Advancing Research and Creativity Award.

References

Jonathan Adler. *Belief's Own Ethics*. MIT Press, 2002.

Pranav Anand and Adrian Brasoveanu. Modal conccord as modal specification. *Proceedings of Sinn and Bedeutung*, 14:19–36, 2010.

Scott AnderBois. On the exceptional status of reportative evidentials. *Proceedings of SALT*, 24:234–254, 2004. doi: https://doi.org/10.3765/salt.v24i0.2424.

Nicholas Asher. Truth conditional discourse semantics for parentheticals. *Journal of Semantics*, 17:31–50, 2000.

Nicholas Asher and Alex Lascarides. *Logics of Conversation*. Cambridge University Press, 2003.

John Langshaw Austin. *How To Do Things With Words*. Harvard University Press, 1962.

Kent Bach. Applying pragmatics to epistemology. *Philosophical Issues*, 18:68–88, 2008. doi: https://doi.org/10.1111/j.1533-6077.2008.00138.x.

Kent Bach and Robert Harnish. *Linguistic Communication and Speech Acts*. MIT Press, 1979.

Mathew Benton and Peter van Elswyk. Hedged assertion. In Sandy Goldberg, editor, *Oxford Handbook of Assertion*, pages 245–263. Oxford University Press, 2020.

Matthew Benton. Two more for knowledge account of assertion. *Analysis*, 71:684–687, 2011.

Matthew Benton. Assertion, knowledge, and predictions. *Analysis*, 72:102–105, 2012.

Matthew Benton. Expert opinion and second-hand knowledge. *Philosophy and Phenomenological Research*, 92(2):492–508, 2016a.

Matthew Benton. Gricean quality. *Noûs*, 50(4):689–703, 2016b.

Matthew Benton. Knowledge is the norm of assertion. In B. Roeber, J. Turri, M. Steup, and E. Sosa, editors, *Contemporary Debates in Epistemology*. Blackwell, 3rd edition, forthcoming.

Diti Bhadra. The semantics of evidentials in questions. *Journal of Semantics*, 37:367–423, 2020.

Martijn Blaauw. Reinforcing the knowledge account of assertion. *Analysis*, 72:105–108, 2012.

Penelope Brown and Stephen Levinson. *Politeness: Some Universals in Language Usage*. Cambridge University Press, 1987.

Lynn Carlson and Daniel Marcu. *Discourse tagging reference manuel (TR-2001-545)*. USC Information Sciences Institute, 2001.

Gennaro Chierchia. Broaden your views: Implicatures of domain widening and the 'logicality' of language. *Linguistic Inquiry*, 37(4):535–590, 2006.

Jonathan Cohen and Andrew Kehler. On convention and coherence. In Gerhard Preyer, editor, *Beyond semantics and pragmatics*, pages 261–283. Oxford University Press, 2018.

Jonathan Cohen and Andrew Kehler. Conversational eliciture. *Philosophers' Imprint*, 21(12):1–26, 2021.

Samuel Cumming, Gabriel Greenberg, and Rory Kelly. Conventions of viewpoint coherence in filim. *Philosopher's Imprint*, 17(1):1–28, 2017.

Keith DeRose. Assertion, knowledge, and context. *Philosophical Review*, 111:167–203, 2002.

Keith DeRose. *The Case for Contextualism: Knowledge, Skepticism, and Context*, volume 1. Oxford University Press, 2009.

Cian Dorr and John Hawthorne. Embedding epistemic modals. *Mind*, 122(488):867–913, 2012.

Michael Dummett. *Frege: Philosophy of Language*. Duckworth, 1973.

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

Andy Egan, John Hawthorne, and Brian Weatherson. Epistemic modals in context. In G. Preyer and G. Peter, editors, *Contextualism in philosophy: Knowledge, meaning and truth*, pages 131–169. Oxford University Press, 2004.

Martina Faller. The discourse commitments of illocutionary reportatives. *Semantics and Pragmatics*, 12(8):1–46, 2019.

Bruce Fraser. Pragmatic competence: The case of hedging. In Gunter Kaltenbröck, Wiltrud Mihatsch, and Stefan Schneider, editors, *New Approaches to Hedging*, pages 15–34. Emerald, 2010.

Manuel Garcia-Carpintero. Assertion and the semantics of force-markers. In C. Bianchi, editor, *The semantics/pragmatics distinction*. CSLI, 2004.

Manuel Garcia-Carpintero. On the nature of presupposition: A normative speech act account. *Erkenntnis*, 85(2):269–293, 2020.

Anastasia Giannakidou and Alda Mari. Epistemic future and epistemic must: Nonveridicality, evidence, and partial knowledge. In Joanna Blaszczak, Anastasia Giannakidou, Dorota Klimek-Jankowska, and Krzysztof Migdalski, editors, *Mood, Aspect, Modality Revisited: New Answers to Old Questions*. University of Chicago Press, 2016.

Anastasia Giannakidou and Alda Mari. The semantic roots of positive polarity: epistemic modal verbs and adverbs in english, greek, and italian. *Linguistics and Philosophy*, 41:623–664, 2018.

Jonathan Ginzburg. *The Interactive Stance*. Oxford University Press, 2012.

Daniel Goodhue. *Must* ϕ is felicitous only if ϕ is not known. *Semantics and Pragmatics*, 10 (14):1–27, 2017.

Paul Grice. *Studies in the Way of Words*. Harvard University Press, 1989.

B. Grosz and C. Sidner. Attention, intentions, and the structure of discourse. *Computational Linguistics*, 12:175–204, 1986.

Daniel Gutzmann. *Use-conditional meaning*. Oxford University Press, 2015.

Valentine Hacquard. *Aspects of modality*. PhD thesis, Massachusetts Institute of Technology, 2006.

Valentine Hacquard and Alexis Wellwood. Embedding epistemic modals in english: A corpus-based study. *Semantics and Pragmatics*, 5(4):1–29, 2012.

Charles Hamblin. Questions in Montague English. *Foundations of Language*, 10(1):41–53, 1973.

Daniel Harris. Intention recognition as the mechanism of human communication. In *Sensations, Thoughts, Language: Essays in Honor of Brian Loar*, pages 11–37. Routledge, 2019.

Jerry Hobbs. Coherence and coreference. *Cognitive Science*, 3:67–90, 1979.

Jerry Hobbs. On the coherence and structure of discourse. Technical Report 85–37, Center for the Study of Language and Information (CSLI), Stanford, CA, 1985.

Jet Hoek, Hannah Rohde, Jacqueline Evers-Vermeul, and Ted Sanders. Scolding the child who threw the scissors: shaping discourse expetations by restricting references. *Language, Cognition, and Neuroscience*, 36(3): 382–399, 2021.

Joan B. Hooper. On assertive predicates. In John P. Kimball, editor, *Syntax and Semantics 4*, pages 91–123. Academic Press, 1975.

Janneke Huitink. Modal concord. a case study in dutch. *Journal of Semantics*, 29(3): 403–437, 2012.

Julie Hunter. Reports in discourse. *Dialogue and Discourse*, 7(4):1–35, 2016.

Julie Hunter and Nicholas Asher. Shapes of conversation and at-issue content. *Semantics and Linguistics Theory*, 26:1022–1042, 2016.

Frank Jackson. On assertion and indicative conditionals. *The Philosophical Review*, 88(4): 565–589, 1979.

Jacques Jayez and Corinne Rossari. Parentheticals as conventional implicatures. In F. Corbin and H. de Swart, editors, *Handbook of French Semantics*, pages 221–229. Stanford CSLI, 2004.

Elsi Kaiser. Effects of anaphoric dependencies and semantic representations on pronoun resolution. In *Anaphora Processing and Applications*, pages 121–130. Heidelberg: Spring, 2009.

Lauri Kartunnen. Possible and must. In *Syntax and semantics 1*, pages 1–20. New York: Academic Press, 1972.

Andrew Kehler. *Coherence, Reference, and the Theory of Grammar*. Stanford University Press, 2002.

Andrew Kehler and Hannah Rohde. A probabilistic reconciliation of coherence-driven and centering-driven theories of pronoun interpretation. *Theoretical Linguistics*, 39(1-2):1–37, 2013.

Andrew Kehler and Hannah Rohde. Prominence and coherence in a bayesian theory of pronoun resolution. *Journal of Pragmatics*, 154:63–78, 2019.

Andrew Kehler, Laura Kertz, Hannah Rohde, and Jeffrey Elman. Coherence and coreference revisited. *Journal of Semantics*, 25 (1):1–44, 2008.

Christoph Kelp. Assertion: A function first account. *Noûs*, 52(2):411–442, 2018. doi: https://doi.org/10.1111/nous.12153.

Todor Koev. Parentheticality, assertion strength, and polarity. *Linguistics and Philosophy*, 44(1):113–140, 2019.

Natasha Korotkova. Evidential meaning and (not-)at-issueness. *Semantics and Pragmatics*, 13(4):1–24, 2020.

Angelika Kratzer. Modality. In Arnim von Stechow and Dieter Wunderlich, editors, *Semantics: an international handbook of contemporary research*, pages 639–650. De Gruyter., 2001.

Jonathan Kvanvig. Assertion, knowledge, and lotteries. In Patrick Greenough and Duncan Pritchard, editors, *Williamson on Knowledge*. Oxford University Press, 2009.

Jennifer Lackey. Norms of assertion. *Noûs*, 41:594–626, 2007.

George Lakoff. Hedges: A study in meaning criteria and the logic of fuzzy concepts. *Journal of Philosophical Logic*, 2(4):458–508, 1973.

Peter Lasershon. Pragmatic halos. *Language*, 75(3):522–521, 1999.

Daniel Lassiter. *Must*, knowledge, and (in)directness. *Natural Language Semantics*, 24(2):117–163, 2016.

Ernest Lepore and Matthew Stone. *Imagination and Convention: Distinguishing Grammar and Inference in Language*. Oxford University Press, 2015.

Brian Loar. *Mind and Meaning*. Cambridge: Cambridge University Press, 1981.

Willem Mak and Ted Sanders. The role of causality in discourse processing: Effects of expectation and coherence relations. *Language and Cognitive Processes*, 28(9):1414–1437, 2013.

Matthew Mandelkern. What *must* adds. *Linguistics and Philosophy*, 42(3):225–266, 2019.

Matthew Mandelkern and Kevin Dorst. Assertion is weak. *Philosophers' Imprint*, 22 (19):1–20, 2022.

William Mann and Sandra Thompson. Rhetorical structure theory: towards a functional theory of text organization. *Text*, 8(3):243–281, 1988.

Lisa Matthewson. Evidential restrictions on epistemic modals. In Luis Alonso-Ovalle and Paula Menendez-Benito, editors, *Epistemic indefinites*, pages 141–160. Oxford University Press, 2015.

Elin McCready. *Reliability in Pragmatics*. Oxford University Press, 2015.

Rachel McKinnon. How do you know that 'how do you know?' challenges a speaker's knowledge? *Pacific Philosophical Quarterly*, 93:65–83., 2012.

Marie-Christine Meyer. *Ignorance and grammar*. PhD thesis, Massachusetts Institute of Technology, 2013. URL https://dspace.mit.edu/handle/1721.1/84420.

G.E. Moore. A reply to my critics. In Paul Schilpp, editor, *The Philosophy of G.E. Moore*. La Salle: Open Court Press, 3rd edition, 1942.

G.E. Moore. *Commonplace Book: 1919-1953*. London: George Allen and Unwin, 1962.

Marcin Morzycki. Metalinguistic comparison in an alternative semantics for imprecision. *Natural Language Semantics*, 19(1):39–86, 2011.

Sarah Moss. On the semantics and pragmatics of epistemic vocabulary. *Semantics and Pragmatics*, 8(5):1–81, 2015. doi: http://dx.doi.org/10.3765/sp.8.5.

Sarah Moss. *Probabilistic Knowledge*. Oxford University Press, 2018.

Sarah Murray. Varieties of update. *Semantics and Pragmatics*, 7:1–53, 2014.

Sarah Murray. *The Semantics of Evidentials*. Oxford: Oxford University Press, 2017.

Sarah Murray and Will Starr. Force and conversational states. In Daniel Fogal, Daniel Harris, and Matt Moss, editors, *New Work On Speech Acts*. Oxford University Press, 2018.

Peter Pagin. Pragmatic enrichment as coherence raising. *Philosophical Studies*, 168 (1):59–100, 2014.

Anna Papafragou. Epistemic modality and truth conditions. *Lingua*, 116:1688–1702, 2006.

Guillermo Del Pinal and Brandon Waldon. Modals under epistemic tension. *Natural Language Semantics*, 27:135–188, 2019.

Andrew Potter. The rhetorical structure of attribution. In *Proceedings of Discourse Relation Tree-banking and Parsing (DISRPT)*, pages 38–49, Minneapolis, MN, 2019. Association for Computational Linguistics.

Ellen F. Prince, Joel Frader, and Charles Bosk. On hedging in physician-physician discourse. In R. D. Pietro, editor, *Linguistics and the professions*, pages 83–97. Ablex: Norwood, NJ, 1982.

Anna Prokofieva and Julia Hirschberg. Hedging and speaker commitment. In *5th International Workshop on Emotion, Social Signals, Sentiment, and Linked Open Data*, pages 10–13, 2014. URL http://www.cs.columbia.edu/~prokofieva/Prokofieva_Hirschberg_Paper.pdf.

Randolph Quirk, Sidney Greenbaum, Geoffrey Leech, and Jan Svartvik. *A Comprehensive Grammar of the English Language*. New York: Longman, 1985.

Jessica Rett. The semantics of emotive markers and other illocutionary content. *Journal of Semantics*, 38(2):305–340, 2021.

Stephen Reynolds. Testimony, knowledge, and epistemic goals. *Philosophical Studies*, 110:139–161, 2002. doi: https://doi.org/10.1023/A:1020254327114.

Craige Roberts. Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics*, 6:1–69, 1996/2012.

Jonathan Schaffer. Knowledge in the image of assertion. *Philosophical Issues*, 18(1):1–19, 2008.

Stephen Schiffer. Intention-based semantics. *Notre Dame Journal of Formal Logic*, 23(2): 119–156., 1982.

Philippe Schlenker. Maximize presupposition and gricean reasoning. *Natural Language Semantics*, 20(4):391–429, 2012.

John Searle and David Vanderveken. *Foundations of Illocutionary Logic*. Cambridge: Cambridge University Press, 1985.

Mona Simion. Assertion: knowledge is enough. *Synthese*, 193(10):3041–3056, 2016.

Mandy Simons. Observations on embedding verbs, evidentiality, and presupposition. *Lingua*, 117(6):1034–1056, 2007.

Ron Smyth. Grammatical determinants of ambiguous pronoun resolution. *Journal of Psycholinguistic Research*, 23:197–229, 1994.

Todd Snider. At-issuness ≠ anaphoric availability. *Proceedings of the Annual Meeting of the Linguistic Society of American*, 2(39):1–15, 2017.

Robert Stalnaker. Pragmatics. *Synthese*, 22: 272–289, 1970.

William B. Starr. What 'if'? *Philosophers' Imprint*, 14(10):1–27, 2014.

Tamina Stephenson. Judge dependence, epistemic modals, and predicates of personal taste. *Linguistics and Philosophy*, 30(4): 487–525, 2007.

Una Stojnić. One's modus ponens: Modality, coherence and logic. *Philosophy and Phenomenological Research*, forthcoming, 2017.

Una Stojnić. *Context and coherence: The logic and grammar of prominence*. Oxford University Press, 2021.

P.F. Strawson. Intention and convention in speech acts. *Philosophical Review*, 73(4):439–460, 1964.

Jonathan Sutton. Stick to what you know. *Noûs*, 39(3):359–396, 2005.

Eric Swanson. How not to theorize about the language of subjective uncertainty. In Andy Egan and Brian Weatherson, editors, *Epistemic Modality*, pages 249–269. Oxford University Press, 2011.

Kristen Syrett and Todor Koev. Experimental evidence for the truth conditional contribution and shifting informational status of appositives. *Journal of Semantics*, 32(3):525–577, 2015.

Judith Tonhauser. Diagnosing (not-)at-issue content. *Proceedings of Semantics of Under-represented Languages of the Americas*, 6:239–254, 2012.

Judith Tonhauser, David Beaver, Craige Roberts, and Mandy Simon. Toward a taxonomy of projective content. *Language*, 89 (1):66–109, 2013.

Stephen Toulmin. *The uses of argument*. Cambridge University Press, 1958.

John Turri. Prompting challenges. *Analysis*, 70:456–462, 2010. doi: https://doi.org/10.1093/analys/anq027.

John Turri. The express knowledge account of assertion. *Australasian Journal of Philosophy*, 89:37–45, 2011.

John Turri. Knowledge guaranteed. *Noûs*, 47:602–612, 2013.

Peter Unger. *Ignorance: The Case for Skepticism*. Oxford: Clarendon Press, 1975.

J.O. Urmson. Parenthetical verbs. *Mind*, 61 (244):480–496, 1952.

Peter van Elswyk. Representing knowledge. *The Philosophical Review*, 130(1):97–143, 2021. doi: https://doi.org/10.1215/00318108-8699695.

Peter van Elswyk and Matthew Benton. Assertion remains strong. *Philosophical Studies*, 180:27–50, 2023.

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

Kai von Fintel and Anthony S. Gillies. Still going strong. *Natural Language Semantics*, 29 (1):91–113, 2021.

Kai von Fintel and Thony Gillies. An opinionated guide to epistemic modality. *Oxford Studies in Epistemology*, 2:32–62, 2007.

Bonnie Webber, Matthew Stone, Aravind Joshi, and Alistair Knott. Anaphora and discourse structure. *Computational Linguistics*, 29(4):545–587, 2003.

Malte Willer. Dynamics of epistemic modality. *Philosophical Review*, 122(1):45–92, 2013.

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

F. Wolf and E. Gibson. *Coherence in Natural Language: Data Structure and Applications*. MIT Press, 2006.

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