Parentheticality, Assertion Strength, and Polarity

Todor Koev

Sentences with slifting parentheticals (such as *The dean greeted the secretary, Jill said*; Ross 1973) grammaticalize an intriguing interaction between truth-conditional meaning and speech act function. In such sentences, the assertion strength of the slifted clause (the non-parenthetical part of the sentence) is modulated by the parenthetical, which provides evidential support (Urmson 1952; Asher 2000; Rooryck 2001; Jayez and Rossari 2004; Davis et al. 2007; Simons 2007; Murray 2014; Maier and Bary 2015; AnderBois 2016; Hunter 2016). Starting with the idea that assertability comes in degrees (Lewis 1976; Davis et al. 2007), this paper develops a probabilistic update model that captures the role of parentheticality as a language tool for qualifying commitments. A crucial role here is played by the rule of Jeffrey conditionalization (Jeffrey 1990), which factors in the uncertainty of the parenthetical information itself. The model also derives certain effects of parenthetical modification not found in regular embedding constructions, including the fact that slifting parentheticals are limited to creating upward-entailing environments.

Keywords: parentheticals, assertion, evidence, polarity, embedding, probability, Jeffrey conditionalization, common ground update

1   Introduction

When performing an utterance, speakers often convey the source of information for what they say. There are several reasons why information sources play such an important role in language. One reason is that speakers want to be believable and may occasionally back up a claim by advancing the available evidence. Another is that specifying an information source may reduce or even remove responsibility for what has been asserted. This paper investigates how such discourse strategies are reflected in language. The empirical focus is on SLIFTING PARENTHETICALS (Ross 1973), which provide evidence and thus modulate the speaker's commitment to the main claim of the utterance (Urmson 1952; Asher 2000; Jayez and Rossari 2004; Davis et al. 2007; Simons 2007; Murray 2014; Maier and Bary 2015; AnderBois 2016; Hunter 2016). The paper also draws comparisons between slifting sentences and the more familiar kind of embedding constructions.

An example of a sentence with a slifting parenthetical is given below.

(1)   The dean greeted the secretary, Jill said.

An utterance of this sentence would imply, in some weak sense, that the dean greeted the secretary and attribute this information to Jill. The relative strength of this implication is contingent on the information contained in the parenthetical, including the lexical semantics of the attitude predicate

(e.g. *say* naturally leads to a weaker claim than *discover*) and the quality of the source (e.g. Jill might be a more trustworthy source than Mary when it comes to the dean's actions). Additional contextual factors may play a role as well, including how strictly the utterer of the sentence is assumed to adhere to the idealized standard of truth when reporting on what Jill said.

On the face of it, slifting sentences are not much different from regular embedding constructions. The slifting sentence in (1) is synonymous with the embedding counterpart in (2). Both sentences constitute speech reports, although in the former case the reporting component is set off parenthetically while in the latter case it is expressed by the matrix clause.

(2)   Jill said that the dean greeted the secretary.

However, syntactic parentheticality has interpretational effects that are not found in regular embeddings. One such effect is that the discourse status of the slifted or non-parenthetical part of the sentence is fixed. An embedding sentence as in (2) does not require that what is reported be AT-ISSUE.[1] If the question under discussion is about what Jill <u>said</u>, then the reported component will indeed be at-issue.[2] If, however, the question under discussion is about what Jill <u>did</u>, then only the entire proposition will be at-issue. With slifting sentences, there is no such freedom. In this case the SLIFTED CONTENT (corresponding to the non-parenthetical part of the sentence) has to be relevant to the question under discussion, otherwise the discourse is infelicitous. This is illustrated by the following contrast.

(3)   Q: Why is Fred not here?
       A: He has quit his job, the secretary told me.

(4)   Q: What happened after the spacecraft sent back the first pictures?
       A: #There was water on Mars, NASA announced.

Another major interpretational effect is that slifting parentheticals are restricted with regard to their polarity. Jackendoff (1972), Ross (1973), Hooper (1975), and Maier and Bary (2015) observe that the occurrence of negation in the parenthetical component is limited. Generalizing this observation, we can say that slifting parentheticals may create upward-entailing but not downward-entailing environments, as seen in (5a). In contrast, matrix clauses are not subject to such restrictions, as shown by (5b).

(5)   a. Snowden is a Russian spy, $\left\{ \begin{matrix} \text{I think} \\ \text{\#I doubt} \end{matrix} \right\}$.

       b. $\left\{ \begin{matrix} \text{I think} \\ \text{I doubt} \end{matrix} \right\}$ Snowden is a Russian spy.

This paper attempts to derive the unique properties of slifting sentences from one basic assumption. Building on suggestions made in previous literature, I assume that slifting sentences perform

---

[1]There are several closely related but distinct notions of at-issueness circulating in the literature (see Koev 2018 for discussion). Here I broadly adopt the notion of at-issueness as relevance to the question under discussion (Simons et al. 2010; Beaver et al. 2017), which seems to have gained the most traction. See (28) below for a precise formulation.

[2]It is often said that in this case the matrix content is not at-issue but is read "discourse parenthetically" (Urmson 1952; Simons 2007; Hunter 2016). While this is an intuitive way of talking, I will argue in subsection 2.3 that – despite appearances – main clause propositions are always at-issue in the intended sense.

two concomitant assertions as part of their conventional meaning (Urmson 1952; Hooper 1975; Bach 1999; Asher 2000; Simons 2007; Maier and Bary 2015; Hunter 2016). More specifically, I assume that the slifted clause encodes the main assertion while the parenthetical adds secondary, typically evidential information. Hooper (1975) puts this idea succinctly; Simons (2007) suggests the same by talking about the main point of the utterance:

> "[...] the effect of complement preposing [=slifting] is to make the complement proposition [=the slifted proposition] the main assertion of the sentence while reducing the original main clause to parenthetical or secondary status. [...] there are two assertions, or two claims to truth. The first one, from the preposed complement [=the slifted clause], is given more importance in this construction, and the parenthetical assertion is clearly subordinated." (Hooper 1975: 95)

> "From a discourse or usage point of view, it seems quite clear that the main point of an utterance of a slifted sentence will be the content of the slifted (i.e. non-parenthetical) clause, with the syntactically parenthetical clause serving a parenthetical function." (Simons 2007: 1039)

I state this assumption as follows.

(6) DOUBLE ASSERTION HYPOTHESIS
    A slifting sentence makes two assertions as part of its conventional meaning. The slifted clause contributes the main assertion and the parenthetical adds secondary information.

This hypothesis entails that (7), repeated from (1), asserts both (7a) and (7b).

(7) The dean greeted the secretary, Jill said.
    a. SLIFTED/MAIN ASSERTION: The dean greeted the secretary.
    b. PARENTHETICAL/SECONDARY ASSERTION: Jill said that the dean greeted the secretary.

One might have qualms about using the label "assertion" here and instead choose to think of (7a) and (7b) as implications of some sort that fall short of reaching the status of assertions. In the case of (7a), the reason is that the slifted component is often too weak to meet the contextual threshold for independent assertions. In the case of (7b), the reason is that the parenthetical component need not be relevant to the question under discussion. Nevertheless, I prefer to view these two components as asserted and investigate how they deviate from other, more standard forms of assertion. The emerging typology of assertion is summarized in the Conclusion.

Although articulated in previous literature, the Double Assertion Hypothesis has barely been put to theoretical use. Here I show how it can be leveraged to develop a comprehensive analysis of slifting sentences and to highlight their core differences to embedding constructions, thus improving on previous treatments of parentheticality and/or evidentiality (see Asher 2000; Jayez and Rossari 2004; Davis et al. 2007; McCready and Ogata 2007; Murray 2014; Maier and Bary 2015; McCready 2015; AnderBois 2016; Hunter 2016; a.o.). The contributions of this paper are several. The paper explicates the dependent nature of slifted assertions, thus demonstrating the relevance of probabilistic updating for semantic theory and emphasizing the usefulness of Jeffrey conditionalization (Jeffrey 1990) for modeling uncertain evidence. The paper also offers an explanation for

the upward monotonicity of slifting parentheticals by paying close attention to their role as utterance modifiers. Overall, we arrive at a layered picture of assertion, where certain linguistic forms may be marked as lacking some of the canonical ingredients of assertive force. This entails that compositional semantics and illocutionary force cannot be cleanly separated; assertion strength may depend on descriptive content, thus blurring the alleged boundary.

Before closing this introductory section, I briefly engage with matters of syntax. It is important to stress that the core structure of slifting sentences turns out to corroborate the proposed analysis. The basic observation is that the view of slifted clauses as contributing the main assertion correlates with their apparent status as main clauses. I list two pieces of evidence in support of this parallelism. One cue comes from internal syntactic properties. Grimshaw (2011: 4–6) presents a battery of arguments in support of the claim that slifted clauses are indistinguishable in form from main clauses. Grimshaw shows among other things that – unlike embedded complements and like main clauses – slifted clauses have to be finite (8), cannot be headed by a complementizer (9), and undergo subject-auxiliary inversion in interrogative contexts (10).[3]

(8) a. *To leave, I promised them.
     b. I promised them to leave.
     c. *To leave.

(9) a. *That it was raining hard, they thought.
     b. They thought that it was raining hard.
     c. *That it was raining hard.

(10) a. Had she made a mistake, he wondered.
     b. *He wondered (whether) had she made a mistake.
     c. Had she made a mistake?

This strongly recommends an analysis of slifted clauses as main clauses (Jackendoff 1972; Potts 2005) over one where they are base-generated as complements and then raised or "s(entence-)lifted" to their surface position (Ross 1973; Rooryck 2001).[4]

A second piece of evidence comes from the observation that slifting sentences cannot be embedded.[5]

(11) a. #Martin now realizes that Sheila is, I had claimed, a luscious yummy.

(Ross 1973: 152)[6]

---

[3]The following data are presented in a slightly adapted form.

[4]Ross (1973) presents several arguments in support of the raising analysis, including the fact that slifting parentheticals license sequence of tense, *de re/de dicto* ambiguities, and strict/sloppy identity readings under VP ellipsis. It seems to me that these primarily semantic phenomena can also be explained under a non-raising analysis by pointing out that the slifted clause, though asserted, supplies the logical argument for the parenthetical and thus enters into a scope relation with it.

[5]The data in (11) also establish that the scope of slifting parentheticals is clause-bounded, in the sense that these expressions necessarily modify the host clause. If this were not the case, the sentences would be acceptable under a reading where the parenthetical modifies the entire structure.

[6]In borrowed examples, I take the liberty of substituting the asterisk (*) with the pound sign (#) whenever the unacceptability is more likely due to pragmatic than syntactic factors.

b. #Selena thinks that Justin, the vocal coach said, is a talented singer.

This property follows if we adopt the view that slifting parentheticals modify the main clause. On the other hand, it is difficult to make sense of if the slifted clause is subordinated, as in this case it remains unclear what prevents the entire construction from being further embedded. While a precise syntactic analysis is not one of the goals here, given the weight of the empirical evidence I will assume that slifted clauses have main clause status and will treat slifting parentheticals as a kind of adverbial expressions.

The paper is structured as follows. Section 2 develops a model that captures the role of slifting parentheticals as utterance modifiers and Section 3 offers an explanation for the restrictions on their polarity. Section 4 is the conclusion.

2 Modulating assertion strength

2.1 Slifted claims as dependent assertions

As already mentioned in the Introduction, the assertion strength of slifted clauses depends on the quality of the evidence contained in the parenthetical. For example, assuming that the reliability of the speaker is kept constant, the utterance in (12a) will likely result in a weaker claim than will the utterance in (12b), which makes a very strong claim. In other words, while (12b) strongly commits the speaker to the slifted claim, (12a) need not do so.[7]

(12) a. Frank is married to a nurse, Lisa claimed.
      b. Frank is married to a nurse, Lisa discovered.

If we schematically represent slifting sentences as C, E (where C stands for "claim" and E stands for "evidence"), the above generalization can be summarized as follows.

(13) STRENGTH DEPENDENCE
      When C, E is uttered, the assertion strength of C is contingent on E.

How can we model dependent assertions like these? A first step toward answering this question involves recognizing the fact that assertability comes in degrees. According to Lewis (1976),

---

[7]As additional support for these intuitions, I presented the following choice to ten American English speakers on Amazon Mechanical Turk.

(i)
> *Imagine that your friend Chris, whom you fully trust, utters one of the following sentences about Jack and Jill, two people you have never met.*
>
> 1. Jack is a bad father, Jill claimed.   2. Jack is a bad father, Jill discovered.
>
> *Which of these sentences makes it more likely that Jack is indeed a bad father?*
>
> [First sentence]   [Second sentence]   [No difference]

Since here the speaker is to be fully trusted and the factual matter is unknown, any contrast in strength must come from the parenthetical. Not surprisingly, nine of the participants picked the second sentence as making a stronger claim (one participant found no difference).

an assertion is warranted when the speaker's degree of confidence in what is being asserted is sufficiently close to full certainty. While Lewis understood the certainty involved in assertion as subjective, or as expressing degrees of belief, I prefer to think of it as objective, or as certainty based on evidence.[8] This choice is motivated by the fact that assertions add propositions to a public body of information and pure belief, no matter how strong, may not be enough or even required. On the other hand, it is quite natural to think of assertion as based on strong but often imperfect evidence, so long as agents are aware of what the contextual standards are. In fact, if we are honest with ourselves, we will admit that most of our everyday assertions fit this exact pattern.

How does this graded view fit into the overall theory of assertion? Assertion is a complex social act that is standardly assumed to have the following components: (i) it imposes some sort of norm or commitment on the speaker, (ii) it constitutes an instruction for the addressee to add the asserted proposition to the common ground, and (iii) it contributes information that is relevant to the question under discussion. It is the first component that is at issue here. Existing theories differ with respect to the type and degree of commitment involved, with some of the options being belief backed by evidence (Searle 1969; Stalnaker 1978; Bach and Harnish 1979; Douven 2006; Hindriks 2007; Lackey 2007), knowledge (Slote 1975; Unger 1975; DeRose 2004; Williamson 2004), or truth (Stenius 1967; MacFarlane 2005; Weiner 2013).[9] The first view (which requires belief backed by evidence) is compatible with the idea of assertability as sufficiently high objective certainty. The second and the third views (which impose knowledge or truth norms) require maximal objective certainty, which seems to set the bar too high and leads to an unrealistic model of communication. Thus, I will assume that the objective certainty of assertion must meet some high contextual threshold that typically lies below full certainty. On top of providing a realistic model of communication, this view allows for thresholds to be manipulated by linguistic means, a feature that will be exploited in the analysis of slifting sentences proposed below.

The idea that assertability comes in degrees has already been invoked in work on parentheticality and evidentiality. Jayez and Rossari (2004) and Davis et al. (2007) distinguish between two language strategies for qualifying a claim. The first strategy involves embedding under a modal operator, which weakens the original claim by manipulating its truth conditions. Under the second strategy, the speaker qualifies the claim by modifying the assertability threshold. Given the evidential nature of slifting parentheticals, slifting sentences arguably lexicalize this latter strategy. Davis et al. (2007: 81) flesh out this strategy in the following three-step procedure, here adapted to slifting sentences.

(14)   If a parenthetical sentence C, E is uttered by an agent a against a threshold θ, then:

    i. a assumes a commitment to having E-type evidence for C,

    ii. θ is readjusted to the reliability of E, and then

    iii. a asserts C against θ.

Notice that the threshold manipulated by the evidence could fall under the default value, and as a result a slifted claim may carry a lower degree of certainty than would a regular assertion in the

---

[8]The full quote reads: "The truthful speaker wants not to assert falsehoods, wherefore he is willing to assert only what he takes to be very probably true. He deems it permissible to assert that A only if $P(A)$ is sufficiently close to 1, where P is the probability function that represents his system of degrees of belief at the time. Assertability goes by subjective probability." (Lewis 1976: 297)

[9]See Brown and Cappelen (2004), Weiner (2007), and MacFarlane (2011) for overviews.

same context. However, if this type of threshold manipulation is assumed to be part of the conventional meaning of parenthetical sentences, there is no danger of confusing dependent assertions like these with a run-of-the-mill high-certainty assertion. We can hypothesize that weak assertions are only allowed if they are explicitly marked as such.[10]

Davis et al.'s main goal is to provide a comprehensive framework for the analysis of evidential markers as illocutionary force operators. It is unclear whether such an essentially pragmatic framework can serve this purpose though, given that in many languages evidential morphemes occur in subordinate clauses (Aikhenvald 2004: 8.1.3). On the other hand, if properly spelled out, this procedure can greatly benefit the analysis of slifting parentheticals, given their evidential flavor and their apparent root clause position (recall the discussion around example (11)).[11] Importantly, the proposed procedure allows us to model the graded strength of slifted clauses without compromising the idea that they are asserted. Unlike standard discourse models, where common ground propositions enjoy full certainty, here we can have contexts that record reliability judgments, including judgments about the strength of slifted claims.

One intuitive way of building uncertainty into our discourse model is to view contexts as associated with probability measures.[12] Specifically, I will assume that contexts minimally consist of a common ground, a probability measure, and an assertability threshold.[13] Common grounds record accepted information and are modeled as sets of propositions. Probability measures assign to every common ground proposition a level of (minimal) objective certainty. As more propositions are added to the common ground, the contextual measure is updated to reflect the gain in information. Finally, thresholds impose a lower bound on the objective certainty of propositions that are added to the common ground.

Here is how probabilistic updating is modeled, doing justice to the intuition that newly asserted information may change the certainty of discourse-old information. Imagine that we are given a prior probability for p (by virtue of p being in the common ground) and that we have just learned q (by virtue of q having been asserted). How should we update the probability of p in light of the new evidence q? A standard mechanism is that of CONDITIONALIZATION, which involves restricting the sample space to the possible worlds compatible with the new evidence and redistributing probabilities within that set. In its simple form, the measure μ conditionalized on the proposition q is defined as follows, for any proposition p.

(15)   CONDITIONALIZATION (simple)

---

[10]To be sure, slifting is not the only way to explicitly mark weak assertions. An anonymous reviewer provides the following example, where the second sentence seems to lower the threshold against which the first sentence is evaluated.

(i)   The burglars fled after the neighbor started to scream. At least, that's what John told me.

I take the fact that the threshold-manipulation mechanism extends beyond slifting to provide additional support for the plausibility of my proposal.

[11]In fact, Davis et al. (2007: 85) themselves suggest this type of extension.

[12]A PROBABILITY MEASURE is a function μ from propositions (sets of possible worlds) to real numbers between 0 and 1 such that (i) μ maps the universal set of worlds to 1, and (ii) μ maps the union of any two disjoint propositions p and q to the sum of $\mu(p)$ and $\mu(q)$. Probability measures have been widely studied and have well-understood properties, although other tools for modeling uncertainty are available as well (see Halpern 2003 and Yalcin 2010 for overviews).

[13]In subsection 2.3, we will add questions under discussion as another parameter.

$$\mu_q(p) = \frac{\mu(p \cap q)}{\mu(q)}$$

However, simple conditionalization is too strict for our threshold-based model, as it presupposes that we always conditionalize on <u>certain</u> information. But what if the evidence itself is uncertain, e.g. because it has been asserted against a non-maximal threshold? In order to model uncertain evidence, work in philosophy and cognitive science has adopted a more general updating mechanism, called JEFFREY CONDITIONALIZATION (Jeffrey 1990: ch.11). It gives us the required gradability by assigning weights to the case that the evidence is true and the case that the evidence is false.[14]

(16)   JEFFREY CONDITIONALIZATION
       $\mu_{q,\theta}(p) = \theta\,\mu_q(p) + [1 - \theta]\,\mu_{\neg q}(p)$

In words, assume that q is potential evidence for p and that q has been asserted against a threshold of θ. The updated probability of p is computed by performing simple conditionalization on q and on $\neg q$, weighting these options out by θ and $1 - \theta$ (respectively), and summing up. This rule then makes the strength of a claim dependent on two things: how strongly the evidence supports the claim in general (because of simple conditionalization) and how certain the evidence itself is (because the support is weighted).

Let me list a few facts about this rule that will be important in the discussion to follow. The first fact is that when the evidence is certain (i.e. $\theta = 1$), Jeffrey conditionalization reduces to simple conditionalization (17a). This means that the latter, more specific rule is just the limiting case of the former, more general rule. The second fact is that conditionalizing on the evidence at the level already assigned to it by the measure does not change the measure (17b). The third fact is that while simple conditionalization on the same proposition returns full certainty ($\mu_q(q) = 1$, as a direct consequence of the definition in (15)), Jeffrey conditionalization on the same proposition at some threshold returns the threshold itself (17c).[15]

(17)   For any proposition q and threshold μ the following facts hold:

       a. $\mu_{q,1} = \mu_q$
       b. $\mu_{q,\mu(q)} = \mu$
       c. $\mu_{q,\theta}(q) = \theta$

The following example provides an illustration of the usefulness of Jeffrey conditionalization when reasoning about uncertain evidence. Let p stand for the proposition "It is raining", q stand for the proposition "Mary said it is raining", and $\mu(p) = 0.5$, $\mu(q) = 0.5$, $\mu_q(p) = 0.7$. That is, we are completely ignorant as to whether p or q are true but we know that learning q would substantially increase the likelihood of p (e.g. because Mary is a fairly reliable source of weather information). What is the posterior probability of p, given that q is uttered at a 0.9 threshold?

- By Bayes' rule, the probability of the evidence q given p is:
$$\mu_p(q) = \frac{\mu_q(p)\,\mu(q)}{\mu(p)} = \frac{0.7 \times 0.5}{0.5} = 0.7.$$

---

[14]This rule falls out from a few simple facts of probability theory. See Appendix A for a proof.
[15]All three facts are derived by direct computation, as shown in Appendix A.

- By the complement rule, the probability of $\neg q$ is:
$\mu(\neg q) = 1 - \mu(q) = 1 - 0.5 = 0.5$.

- By Bayes' rule and the complement rule, the probability of p given $\neg q$ is:
$\mu_{\neg q}(p) = \dfrac{\mu_p(\neg q)\,\mu(p)}{\mu(\neg q)} = \dfrac{[1 - \mu_p(q)]\,\mu(p)}{1 - \mu(q)} = \dfrac{[1 - 0.7] \times 0.5}{1 - 0.5} = 0.3$.

- Finally, by Jeffrey conditionalization and given that q is uttered at a threshold of 0.9, the posterior probability of p becomes:
$\mu_{q,0.9}(p) = 0.9\,\mu_q(p) + [1 - 0.9]\,\mu_{\neg q}(p) = 0.9 \times 0.7 + 0.1 \times 0.3 = 0.66$.

The probability of rain thus increases from 50% to 66%. This is lower than the 70% produced by simple conditionalization, since this latter rule does not take into account that we are only 90% certain that Mary said it was raining.

2.2 Update rules

We can now be more specific about the semantics of sentences with slifting parentheticals. I start off by providing a general update rule for non-parenthetical sentences S. Below, CG is a common ground, μ is a probability measure, θ and τ are threshold values, p is a proposition, and $+$ is the update function.

(18) GENERAL UPDATE
$\langle CG, \mu \rangle +_\theta S = \langle CG \cup \{[\![S]\!]_\theta\}, \mu_{[\![S]\!],\theta} \rangle$,
provided that for all $p_\tau \in CG : \tau \leq \mu_{[\![S]\!],\theta}(p)$

This rule states that a non-parenthetical assertion has two effects on the context: it adds the asserted proposition (indexed by the contextual threshold) to the common ground and it Jeffrey-conditionalizes the measure on that proposition at the threshold. As a result, the common ground propositions are reassigned probabilities that reflect both the content and the certainty of the newly added proposition. The definedness condition ensures that these probabilities are at least as high as the respective thresholds at which the propositions have been originally asserted. This excludes a case where conditionalization on a newly uttered proposition (at the given threshold) makes a common ground proposition too uncertain to have been asserted in the first place. We thus prevent incompatible information from entering the common ground, although here the degree of incompatibility is quantified by the thresholds.

Notice that the contextual measure is conservative in the following sense. An asserted proposition conditionalizes the measure at the given threshold even when the confidence in that proposition exceeds the threshold. This means that the contextual measure assigns to every common ground proposition the <u>minimal</u> degree of certainty compatible with the available information. But this does not mean that the model fails to represent our ignorance about what the precise probabilities are. On the one hand, we can encode this ignorance directly by letting conditionalization operate on <u>sets</u> of measures at levels that meet or exceed the contextual threshold (see also Halpern 2003: 2.3; Moss 2015; Rudin 2018; Bradley 2019). Following this line, we might want to substitute (18) with (19) (the definedness condition is ignored here).

(19) $\langle CG, M \rangle +_\theta S = \langle CG \cup \{[\![S]\!]_\theta\}, \{\mu_{[\![S]\!],n} \mid \mu \in M \text{ and } \theta \leq n \leq 1\} \rangle$

However, this move is not required, because all the necessary information is already contained in our single-measure contexts. Thus, (19) can be derived from (18) by mapping a context $\langle CG, \mu \rangle$ to a context $\langle CG, M \rangle$, where M is the set of measures that assign to any given proposition in CG a value that is at least as high as the value assigned to that proposition by μ. This is achieved by the following "raise" ($\uparrow$) operation.

(20) $\uparrow \langle CG, \mu \rangle = \langle CG, \{ \mu' \mid \text{for all } p \in CG : \mu(p) \leq \mu'(p) \} \rangle$

The moral of the story is this: We should keep in mind that the measures of our contexts provide lower bounds rather than precise values for common ground propositions. However, since the upper bound for asserted propositions is always full certainty, our ignorance about the precise values can be retrieved without complicating the semantics.[16]

I now discuss the interpretation of slifting sentences. I propose that slifting sentences are subject to the following rule, which spells out two consecutive updates, in line with the Double Assertion Hypothesis. The global context is first updated with the parenthetical proposition at the default threshold, and the resulting context is further updated with the slifted claim at the certainty level assigned to that claim by the newly updated measure.

(21) PARENTHETICAL UPDATE
$\langle CG, \mu \rangle +_\theta C, E = \langle CG, \mu \rangle +_\theta E(C) +_{\mu'(\llbracket C \rrbracket)} C,$
where $\mu' = \mu_{\llbracket E(C) \rrbracket, \theta}$

This rule may look complicated, but it boils down to something rather simple, as the following derivation demonstrates.[17]

(22) $\langle CG, \mu \rangle +_\theta C, E = \langle CG, \mu \rangle +_\theta E(C) +_{\mu'(\llbracket C \rrbracket)} C$ by (21)
$= \langle CG \cup \{\llbracket E(C) \rrbracket\}, \mu' \rangle +_{\mu'(\llbracket C \rrbracket)} C$ by (18)
$= \langle CG \cup \{\llbracket E(C) \rrbracket\} \cup \{\llbracket C \rrbracket\}, \mu'_{\llbracket C \rrbracket, \mu'(\llbracket C \rrbracket)} \rangle$ by (18)
$= \langle CG \cup \{\llbracket E(C) \rrbracket, \llbracket C \rrbracket\}, \mu'_{\llbracket C \rrbracket, \mu'(\llbracket C \rrbracket)} \rangle$ by set union
$= \langle CG \cup \{\llbracket E(C) \rrbracket, \llbracket C \rrbracket\}, \mu' \rangle$ by fact (17b)
$= \langle CG \cup \{\llbracket E(C) \rrbracket, \llbracket C \rrbracket\}, \mu_{\llbracket E(C) \rrbracket, \theta} \rangle$ by (21)

The last line in (22) restates the rule in (21) in more transparent terms. It says that when a slifting sentence C, E is uttered, the common ground is enriched with the evidential proposition $\llbracket E(C) \rrbracket$ and the claim $\llbracket C \rrbracket$, while the measure function μ is conditionalized on $\llbracket E(C) \rrbracket$ at the contextual threshold θ. Even though (21) indicates further conditionalization on $\llbracket C \rrbracket$ at a threshold of $\mu_{\llbracket E(C) \rrbracket, \theta}(\llbracket C \rrbracket)$, this move is rendered vacuous by the fact stated in (17b). This is arguably a good result. Since slifted claims wear their evidence on their parenthetical sleeve, so to speak, they automatically match the threshold resulting from the evidence.

---

[16]An anonymous reviewer asks whether the semantics predicts that asserting a proposition at some threshold should entail a simultaneous assertion of the negated proposition at one minus the threshold. Now, it is true that if a sentence S is asserted at θ, the updated measure $\mu_{\llbracket S \rrbracket, \theta}$ will assign θ to $\llbracket S \rrbracket$ (by the fact in (17c)), and hence it will also assign $1 - \theta$ to $\llbracket \neg S \rrbracket$ (by the complement rule of probability theory). Despite that, the semantics does not predict a concomitant assertion of $\neg S$, for the simple reason that only $\llbracket S \rrbracket$ will be added to the common ground. If the negated proposition were allowed to enter the common ground as well, we would expect all kinds of unattested linguistic effects, including that this proposition should be anaphorically accessible in subsequent discourse. That this is not so can be seen from the fact that the second sentence in *It's raining. John told me that* cannot mean "John told me it's not raining".

[17]From now on, I will omit indexing common ground propositions by thresholds.

The update rule for slifting sentences correctly captures our intuitions about strength. To see how, consider a context in which *It's raining, John said* has been uttered. The strength of the parenthetical assertion *John said it's raining* is predicted to be at least $\mu_{[\![\text{John say(rain)}]\!],\theta}([\![\text{John say(rain)}]\!])$, which by the fact in (17c) amounts to θ, the default threshold of the context. This is why such assertions are perceived as strong. In turn, the strength of the main assertion *It's raining* must meet the value $\mu_{[\![\text{John say(rain)}]\!],\theta}([\![\text{rain}]\!])$. This simply means that the likelihood of rain will depend on the quality (and certainty) of the information that John said it is raining. We have derived the Strength Dependence generalization in a straightforward manner.

2.3 Hedging, strength, and relevance

The analysis of slifting sentences given above predicts that slifted claims should be able to be hedged to the degree to which embedded claims can. The reason is that the resulting measures are equivalent and would assign to either claim the same (minimal) probability. In (22) above, we computed that the dynamic effect of slifting sentences amounts to (23a). In turn, embedding sentences are non-parenthetical and thus are subject to the rule in (18); their effect on the context must then be as in (23b). The two output contexts derive the same measure and only differ with respect to what is added to the common ground.

(23) a. $\langle CG, \mu \rangle +_\theta C, E = \langle CG \cup \{[\![E(C)]\!], [\![C]\!]\}, \mu_{[\![E(C)]\!],\theta} \rangle$
      b. $\langle CG, \mu \rangle +_\theta E(C) = \langle CG \cup \{[\![E(C)]\!]\}, \mu_{[\![E(C)]\!],\theta} \rangle$

In spite of this, Jackendoff (1972), Asher (2000), and Murray (2014) draw a sharp empirical distinction between denials of slifted and embedded claims. On the face of it, the data below could be taken as evidence that slifted clauses make a stronger claim than do embedded clauses, something that would be unexpected under the current proposal.

(24) (Jackendoff 1972: 97)
      a. Myrtle will come tomorrow, Margaret believes, (#but she actually came yesterday).
      b. Margaret believes that Myrtle will come tomorrow, but she actually came yesterday.

(25) (Asher 2000: 36)
      a. #John, Mary assures us, can be trusted, but I don't trust him.
      b. Mary assures us that John can be trusted, but I don't trust him.

The status of this contrast is less clear than stated, though, as some of the English speakers I consulted expressed hesitance in their judgments. It is also important to ask whether the judgments change if the follow-up clause does not categorically deny the claim but leaves open the possibility that the slifted clause is true. In order to provide robust and detailed data on these sensitive intuitions, I conducted an experimental study, which is described below.

The experiment was an acceptability task aimed at revealing the extent to which slifted and embedded claims can be hedged by the speaker in conjoined follow-ups. The experiment employed a 2×2 design, and included as predictors Sentence (Parenthetical v. Embedding) and Follow-up (Strong, e.g. *but she is wrong*, v. Weak, e.g. *but she might be wrong*). Eight sets of stimuli were

created, where the four sentences in each set minimally differed across the four experimental conditions, thus producing a total of 32 tokens.[18] To these tokens were added eight filler items, which also served as controls. Responses were recorded on a three-point Likert scale. A sample trial is given below, where Sentence=Parenthetical and Follow-up=Strong.

(26)

> *In your opinion, is the following sentence acceptable?*
>
> Lucy bought a new car, my brother said, but she didn't.
>
> [Yes]   [Somewhat]   [No]

40 participants, all native speakers of American English, were recruited through the Prolific crowdsourcing platform. Four participants were removed because they failed more than one of the controls, thus bringing the number of participants down to 36. Each participant saw all fillers and two tokens per experimental condition.

The obtained results are summarized in Table 1. Looking at the stacked barplot and the means, we see that slifted claims are generally harder to deny or doubt than embedded claims, irrespective of the shape of the follow-up.[19] This was confirmed by the statistical analysis. An ordinal logistic regression model revealed that sentence type is a highly significant factor for acceptability ($\beta = 2.072, SE = 0.388, t = 5.342, p < 0.001$). In turn, follow-up type was not found to be a contributing factor ($\beta = 0.447, SE = 0.419, t = 1.068, p = 0.286$).[20]

We can draw two major conclusions from these findings. The first conclusion is that slifted claims can in principle be hedged: Participants found such examples to be acceptable in well over a third of the cases (38.2% on average). This suggests that speakers can be committed to a slifted claim to various degrees, including not being committed at all. Despite this, a second conclusion is that slifted claims are significantly more difficult to hedge than embedded ones. Overall, we find that while the claimed contrast in (24)-(25) is real, it is not a categorical one but rather a matter of degree.

Do these conclusions challenge the prediction that slifted and embedded claims are of equal strength? I argue that they do not. In fact, a pattern such as the one above is to be expected if we notice that slifted and embedded claims have different discourse statuses. I will take this difference to be a confounding factor that makes hedging of slifted claims less natural.

Recall from (3)-(4) above that slifted clauses have to be at-issue or relevant to the question under discussion, otherwise the discourse is ill-formed. I will assume that main clause propositions are at-issue across contexts and that slifted clauses inherit this property by virtue of their main clause status. Empirical arguments for the main clause status of slifted clauses were presented at the end of Section 1. To motivate the at-issueness requirement on main clause propositions, consider the following contrast.

(27)   Q: What's the weather going to be tomorrow?

---

[18]The test items included parenthetical or matrix clauses with verbs of saying or doxastics. Factive verbs were excluded because they entail the claim and disallow hedging overall. Appendix B contains the full list of stimuli.

[19]The means were computed by coding "acceptable" as 1, "somewhat acceptable" as 0.5, and "unacceptable" as 0.

[20]The statistical analysis was performed in the R statistical software and used the polr() function, available in the MASS package.

Relative Acceptability
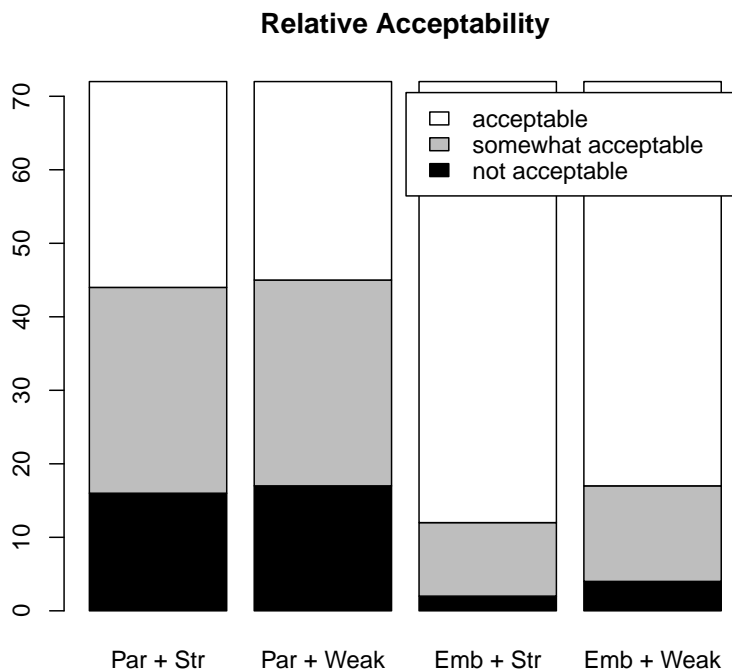
condition	acceptable	somewhat acceptable	not acceptable	mean
Par + Str	28 (38.9%)	28 (38.9%)	16 (22.2%)	0.58
Par + Weak	27 (37.5%)	28 (38.9%)	17 (23.6%)	0.57
Emb + Str	60 (83.3%)	10 (13.9%)	2 (2.8%)	0.90
Emb + Weak	55 (76.4%)	13 (18.1%)	4 (5.6%)	0.85

Table 1: Relative acceptability of hedged slifted and embedded claims.

$$\text{A: Mary} \left\{ \begin{array}{c} \text{said} \\ \text{\#dreamed} \end{array} \right\} \text{that it's going to rain.}$$

The first version of the answer is felicitous. However, this is not because the embedded complement is relevant to the question under discussion but rather because a speech report would typically change the likelihood of its complement being true. If all the theory of at-issueness required were that some part of the utterance be relevant, the second version of the answer would be felicitous as well, but it is not. Intuitively, the second version is odd because – magical scenarios aside – dream reports have no bearing on the likelihood of their complements and thus the sentence as a whole says nothing about the question at hand. We can conclude that in a felicitous discourse main clause content is always relevant to the question under discussion. Embedded complements may or may not be relevant; we need not impose any restrictions here.

The notion of relevance that I will endorse generalizes the stricter notion of Groenendijk and Stokhof (1984), Simons et al. (2010), a.o. and is commonly adopted in the literature. According to this stricter notion, a relevant proposition provides a (full or partial) answer to the question, i.e. it decides on at least one question alternative. However, this definition fails to capture the intuition that the first version of the answer in (27) is relevant, even though it does not decide on

any question alternative. As suggested in Büring (2003) and Simons et al. (2010), we need a more sensitive notion according to which it is enough that a relevant proposition, if assumed to hold true, modifies the likelihood of some question alternative. The idea can be spelled out as follows.

(28) RELEVANCE
      A proposition p is RELEVANT to a question Q with respect to a measure μ iff for some $q \in Q : \mu_p(q) \neq \mu(q)$.

This notion of relevance subsumes the stricter notion mentioned above in the following sense. A complete answer amounts to selecting a unique question alternative by assigning it a probability of 1, while excluding all remaining alternatives by assigning them a probability of 0. A partial answer amounts to assigning to at least one alternative a probability of 0 and bumping up the likelihoods of the remaining alternatives in order to fill up the probability space. Finally, a "suggested" answer as in (27A) redistributes weights among alternatives without assigning to any of them a probability of 0.

    We can now impose relevance constraints on sentences with different shapes in the form of definedness conditions on updates. We require that main clause content be relevant, or else the update rule cannot apply.

(29) a. $\langle CG, \mu \rangle +_{\theta,Q} S = \langle CG, \mu \rangle +_\theta S,$
          provided that $[\![S]\!]$ is relevant to Q with respect to μ
      b. $\langle CG, \mu \rangle +_{\theta,Q} C, E = \langle CG, \mu \rangle +_\theta C, E,$
          provided that $[\![C]\!]$ is relevant to Q with respect to μ

These relevance constraints can explain the reduced acceptability of hedged slifted claims as follows. According to (29), slifted but not embedded claims have to be relevant to the question under discussion. In the case of hedged slifted claims, it would then be odd for the speaker to mark a proposition as relevant and thus intended to be added to the common ground, and then immediately go on to doubt or disavow it. This is especially true in view of the availability of a competing embedding construction, which does not mark the embedded claim as relevant and more naturally allows for hedging. If the attested contrast in hedging possibilities is indeed due to information-structural factors, we can understand why this contrast is not very sharp. We can also uphold the view that slifted and embedded claims do not intrinsically differ in strength.[21]

2.4 Previous views on slifted claims

A lot of recent works in semantics and pragmatics have offered detailed accounts of phenomena like evidentiality, parentheticality, hedging, sentence adverbials, attitude reports, etc., and a subset of these works make predictions about the status and strength of slifted claims (Asher 2000; Jayez and Rossari 2004; Davis et al. 2007; McCready and Ogata 2007; Murray 2014; Maier and Bary 2015; McCready 2015; AnderBois 2016; Hunter 2016; a.o.). This subsection critically evaluates the different predictions and points out the empirical problems that most of them face.

---

[21]A second experimental task juxtaposed slifting and embedding sentences and asked participants to determine which sentence, if any, makes the claim more likely to be true. This experiment rendered no significant results, so it is not reported in detail here. Notice, though, that the lack of a statistically significant contrast is fully consistent with the assumption that slifted and embedded claims are of equal strength.

Asher (2000) and Hunter (2016) frame their accounts of parenthetical expressions in terms of Segmented DRT, which views discourse segments as connected into a coherent whole (Asher and Lascarides 2003). According to these accounts, slifted claims (in reportative contexts) are linked to the parenthetical proposition through a veridical rhetorical relation such as Evidence or Source. Since veridical relations entail their arguments, it follows that speakers are fully committed to slifted claims. However, this prediction is belied by the experimental task presented in the previous subsection, which demonstrated that slifted claims can be hedged. Hunter adds the caveat that the commitment requirement can be waived if the slifting sentence extends a previous report, as in the corpus example below. (The initial and the dependent speech contexts are underlined.)

(30) An association spokeswoman, Elinore Boeke, said travelers should check with public health officials about destinations they planned to visit and cited the C.D.C.'s current travel advisories, which suggest only that all visitors avoid mosquito bites by using repellent and long clothes. Cruise ships publish daily fliers on health and safety and instruct passengers on how to avoid bites, Ms. Boeke added.   (Hunter 2016: 29)

Since the above experiment did not control for discourse factors like these, it may be thought that it is these types of environments that participants had in mind when judging hedged cases as acceptable. However, Hunter's account still predicts that slifted claims in fresh parenthetical reports should be fully committing. But this seems problematic in view of the data in (31), where a first-time parenthetical report is felicitously hedged by the speaker.

(31) a. There we met an Italian named Giuseppe Ugge, whose job, he told us, was to introduce the United Arab Emirates royal family to foreign business executives. He looked like the prototypical con man and what he was telling us was almost certainly not true.
                                                                    (based on a COCA example)

      b. People, my uncle once told me, are more difficult to work with than machines. But I don't think that's true. This is just his opinion.

Both Asher's and Hunter's accounts make predictions that appear to be too strong, at least when applied to a subset of the data.

A more nuanced position is that the speaker is committed not to the slifted claim directly but to a modalized version of it. Different ideas in this ballpark have been defended in the literature. Jayez and Rossari (2004), for example, propose that slifted claims are prefixed by a necessity modal that is anchored to the evidence source. Under their view, *It's raining, Mary said* essentially states that it is raining in all of the worlds in which what Mary says is true. Since this is equivalent in meaning to *Mary said it's raining*, it remains unclear what predictions Jayez and Rossari's account makes about the strength of the slifted claim itself. Another, more explicit choice is that the speaker is committed to the possibility that the slifted proposition is true (Murray 2014). However, as Murray herself notices, some uses entail no commitments to the slifted claim at all. At the same time, recalling the discussion around example (12), a full commitment will ensue if the slift is headed by a factive verb like *discover* or *realize*. Treating slifted claims as being possibly true is then both too strong and too weak. What we need is a mechanism that is flexible enough to capture the varying strength of slifted claims depending on the shape of the parenthetical, as under the current account. Finally, one may try to argue that slifted claims are prefixed by operators with a variable force. In their analysis of Japanese inferential evidentials, McCready and Ogata (2007) introduce

modal operators whose strength is evaluated relative to a probability measure conditionalized on the given evidence. Adapted to slifting sentences, this amounts to the speaker being committed to the modalized slifted claim, where the modal is interpreted relative to a probability measure that is conditionalized on the parenthetical information. This would capture the dependent strength of slifted claims by introducing flexible modal operators rather than by manipulating thresholds. It would be interesting to see how such an account fares relative to the current proposal when properly worked out for parentheticals.

The weakest possible view is that slifted claims entail no commitments whatsoever. Maier and Bary (2015) and AnderBois (2016) go one step in this direction. According to these authors, parenthetical reports entail full commitment to the evidential component but the speaker proposes to add to the common ground a different proposition, i.e. the slifted claim.[22] Since this view is limited to reportative uses of slifting parentheticals (for example, it excludes parentheticals with factive verbs), it does not make predictions about the strength of slifted claims in all possible contexts. Still, it may seem dubious that slifted claims in parenthetical reports <u>never</u> entail commitments. While the experimental data presented in the previous subsection shows that this is possible, it would be strange if this were the only option.

3 Polarity restrictions

3.1 Monotonicity in slifting parentheticals

Other things being equal, one might expect that slifting parentheticals would be able to take the same shapes as matrix clauses. For example, the fragment *Mary said* can be used parenthetically or it can serve as a matrix clause. As it turns out though, slifting parentheticals are more restricted in shape than matrix clauses. Jackendoff (1972), Ross (1973), Hooper (1975), and Maier and Bary (2015) notice that slifting parentheticals need to make a "positive" or "affirmative" import. At least initially, it seems that slifting parentheticals cannot contain negation or lexically negative verbs (32). This is surprising, given that matrix clauses are not subject to this restriction (33).

(32) (Jackendoff 1972: 97)
    a. John is, I think, a fink.
    b. #John is, $\left\{ \begin{array}{c} \text{I don't think} \\ \text{I doubt} \end{array} \right\}$, a fink.

(33) $\left\{ \begin{array}{c} \text{I think} \\ \text{I don't think} \\ \text{I doubt} \end{array} \right\}$ John is a fink.

However, the same authors also observe that there is no syntactic ban per se on slifting parentheticals hosting negation or negative predicates. If a slift contains a lexically negative verb which itself is negated, the sentence is often acceptable.[23]

---

[22] AnderBois (2016) calls such cases "asymmetric" or "decoupled" assertions, because here what is asserted (the slifted claim) diverges from what the speaker is committed to (the parenthetical component).

[23] The following examples are culled from previous literature and some speakers find them degraded. However, notice that the proposal developed in this section rules out non-upward-monotone slifts but makes no predictions as

(34)	John is, I don't doubt, a fink.	(Jackendoff 1972: 97)
(35)	Mushrooms are great on diets, I don't doubt.	(Ross 1973: 155)
(36)	It's a long shot, I don't deny.	(Hooper 1975: 107)

The positive/affirmative import restriction on slifts must be a semantic one and it must constrain the slifting parenthetical as a whole rather than its main predicate. This restriction should then be stated in terms of the monotonicity properties of the entire parenthetical.

The monotonicity properties of operators have been widely discussed in the literature on polarity items. The two crucial notions of upward and downward entailingness are standardly defined as follows (e.g. Ladusaw 1980).

(37)   a. An operator O is UPWARD ENTAILING iff for any propositions p and q, if p entails q then $O(p)$ entails $O(q)$.

      b. An operator O is DOWNWARD ENTAILING iff for any propositions p and q, if p entails q then $O(q)$ entails $O(p)$.

Asher (1987) develops a semantic typology of a wide range of propositional attitude predicates, partially based on their monotonicity properties. Asher lists verbs like *say*, *think*, *believe* as upward entailing and verbs like *deny* or *doubt* as downward entailing.[24] Indeed, if Jimmy believes I own a Porsche (and he also knows Porsches are cars), he must believe that I own a car. In turn, if Jimmy doubts I own a car (and he knows Porsches are cars), he must doubt that I own a Porsche.

The polarity generalization can now be stated as follows.

(38)   UPWARD MONOTONICITY
A slifting sentence S, E can be felicitously uttered only if E expresses an upward-entailing operator.

Assuming that no operator can be both upward and downward entailing, Upward Monotonicity correctly bans slifts from creating downward-entailing environments. It also rules out slifts that express non-monotone operators, i.e. operators that are neither downward nor upward entailing. Do such slifting parentheticals exist? The attitude verb *lie* appears to be non-monotone; for example, there are no entailment relations between (39a) and (39b).[25] Upward Monotonicity then correctly predicts that the sentences in (40), which host non-monotone slifts, are not acceptable.

(39)   a. Jeremy lied that he was born in Vancouver.

      b. Jeremy lied that he was born in Canada.

(40)   a. #Paul had never known his mother, he lied.

      b. #The CEO, his girlfriend lied, is a true gentleman.

---

to whether a particular upward-monotone (or non-monotone) slift is acceptable. There may be further restrictions on slifting that play a role here (see Hooper 1975 and Haddican et al. 2014 for discussion).

[24] I will sometimes loosely speak of words or phrases as having certain monotonicity properties. What is meant is that the operators that those words or phrases denote have those properties. I will also call attitude verbs like *think* upward entailing, meaning that the corresponding operator is upward entailing with respect to its propositional argument.

[25] The reason *lie* is non-monotone seems to be that *a lied that p* implies both *a said that p* and *not p*. Since *say* is upward entailing while *not* is downward entailing, the combined effect is that of non-monotonicity. I am indebted to Emar Maier for suggesting to me that the verb *lie* expresses a non-monotone operator.

There are two cases of seeming violations of Upward Monotonicity that deserve further mention. The first case arises when a speaker echoes a previous utterance and disagrees with it by using a segment that resembles a negated slifting parenthetical. This use is illustrated in (41), which can be read as a self-directed question/answer pair. The second case involves a negated main clause followed by a negated slifting parenthetical whose apparent function is to qualify the categorical denial expressed by the first part of the sentence (42).

(41)   Mushrooms are great on diets, I don't think so.

(42)   Matt doesn't like phonology, I don't think.

It is fairly clear that (41) is not a slifting sentence at all. The obvious thing to notice is that the argument slot of the alleged parenthetical here is filled by a *so*-anaphor and without it the sentence is unacceptable. By contrast, genuine slifts rule out *so*-insertion in this position (cf. **The dean greeted the secretary, Jill said so*). While examples as in (42) do exhibit the grammatical form of slifting, they have quirky distributional and interpretational properties. One important observation, made in Ross (1973) and Hooper (1975), is that in such examples the parenthetical negation is licensed by the main clause negation. Just any downward-entailing operator in the main clause cannot support a negated slift (cf. *#Few people like phonology, I don't think*) and just any downward-entailing slift cannot compose with a negated main clause (cf. *#Matt doesn't like phonology, I doubt*). Another observation is that the missing argument of the parenthetical need not be a clause (cf. *Matt doesn't like phonology, I don't think he does*). Yet another difference to regular slifting sentences is that here the parenthetical has to occur sentence-finally (cf. *#Matt, I don't think, doesn't like phonology*). But perhaps the main giveaway is that the negation in the parenthetical is semantically vacuous: (42) is synonymous with *Matt doesn't like phonology, I think*. I will then assume that negated slifts of this type arise through some sort of negative concord mechanism and are actually upward entailing, as previously suggested in Hooper (1975), Culicover (1992), and Huddleston and Pullum (2002: 845). If so, the data in (41)-(42) do not posit a real challenge to Upward Monotonicity.

3.2   Explaining Upward Monotonicity

Upward Monotonicity is just an empirical generalization. But where does it come from? There are two related ideas that have been floated in the literature. The first idea is that Upward Monotonicity falls out from a constraint on the strength of slifted claims in some absolute sense. Hooper (1975), Scheffler (2009), and Hunter (2016) propose that upward-entailing parentheticals convey a sufficiently high level of confidence in the slifted claim while downward-entailing parentheticals do not. Employing the formalism introduced in Section 2, this constraint can be stated as follows.

(43)   SUFFICIENT STRENGTH
        A slifting sentence C, E can be felicitously uttered in a context with measure μ and threshold θ only if $\tau \leq \mu_{\llbracket E(C) \rrbracket, \theta}(\llbracket C \rrbracket)$, where τ is the minimal threshold required for high certainty.

Another idea is that Upward Monotonicity is tied to the evidential role of slifting parentheticals in a relative sense. Haddican et al. (2014) and Maier and Bary (2015) suggest that upward-entailing

operators provide confirming evidence for the target proposition while downward-entailing operators do not. We can formalize this idea by requiring that the parenthetical information raise the likelihood of the main sentence.

(44) EVIDENTIAL SUPPORT
       A slifting sentence C, E can be felicitously uttered in a context with measure μ and threshold θ only if $\mu([\![C]\!]) < \mu_{[\![E(C)]\!], \theta}([\![C]\!])$.

What Sufficient Strength and Evidential Support have in common is that the posterior probability of the slifted claim is required to meet or exceed some value (the minimal threshold in the former case, the prior probability in the latter case). Thus, both explanations suggest that downward-entailing parentheticals fail to back up the slifted claim to the required degree.[26] I now discuss some of the empirical challenges for these explanations and propose an alternative.

   Consider first the contrast in (45). The non-factive segment *he claims* will generally convey a lower degree of confidence in the slifted claim than will the factive segment *he regrets*, and yet only the former segment can be used parenthetically. Since Sufficient Strength and Evidential Support are stated as necessary (rather than sufficient) conditions on slifting, the unacceptability of (45b) cannot be held against either of these constraints. However, given that (45a) makes a weaker claim and is acceptable, this sentence must obey both constraints. (45b), which results in a stronger claim, must then obey them as well. But this means that (45b) needs to be ruled out by some additional constraint, and thus the contrast in judgment remains unexplained.

(45) a. Noah went to Harvard, he claims.
      b. #Noah went to Harvard, he regrets.

There is also data that directly contradicts the above constraints. These are cases where the parenthetical renders the claim too weak, or even false, but the sentence is acceptable. Imagine that Putin is a rogue leader who typically makes false claims in order to deceive his enemies. A news headline as in (46) is acceptable in this context, even though both constraints should be violated.[27] Sufficient Strength should be violated because if Putin rarely tells the truth, what he says is more likely to be false than it is likely to be true, and this is too low a bar to guarantee high certainty, even under a very tolerant notion of high certainty. Evidential Support should be violated as well, given that Putin's utterance makes the claim less (not more) likely to be true.

(46) There are Russian troops in Ukraine, Putin claims.

The same point can be made even more forcefully by citing good examples where the slifted claim is rendered plain false. If someone is untrustworthy to the point where they <u>never</u> tell the truth, the sentence in (47) should be ruled out by both of the above constraints. And yet, this sentence is judged to be acceptable.

---

[26]Though similar in spirit, notice that Sufficient Strength and Evidential Support are logically independent. The likelihood of a claim may stay above the minimal threshold even if the claim receives negative support from the parenthetical, but the size of the effect is small. Conversely, a claim may get a positive but minor boost from the parenthetical, so that its likelihood remains below the minimal threshold.

[27]A reviewer wonders whether (46) is only acceptable in news headlines, stating that headlines allow for a lot of linguistic behavior that is otherwise not allowed. To me, the important point is that this sentence does have a natural habitat, despite the unreliability of the evidence it carries with it.

(47)   This box is empty, the liar said.

In view of all these data, the upward monotonicity of slifting parentheticals does not seem to follow from constraints on (absolute or relative) strength. I will now argue that it is crucially linked to the assertion status of slifted clauses.

In order to arrive at the explanation I have in mind, it will be useful to shift our focus from the parenthetical to the slifted clause and reframe the problem a bit. Recall that according to the proposed analysis, the assertion of the slifted clause is modified by the parenthetical. If we assume that slifted clauses in particular – or main clauses in general – are positive polarity items (PPIs), we can begin to understand why their modifiers are upward entailing. Thus, (48a) hosts a main clause under an upward-entailing (UE) operator and the sentence is acceptable, while (48b) hosts a main clause under a downward-entailing (DE) operator and the sentence is not acceptable.

(48)   a. [Snowden is a Russian spy]$_{PPI}$, [Mary thinks]$_{UE}$
       b. #[Snowden is a Russian spy]$_{PPI}$, [Mary doubts]$_{DE}$

Why do main clauses behave like positive polarity items? I suggest this is because such clauses are asserted and assertions should allow speech participants to draw entailments, i.e. to derive logically weaker statements from what is said. A general constraint along these lines can be stated as follows.

(49)   ENTAILMENT CLOSURE
       A sentence can be felicitously asserted only if all of its entailments meet the contextual threshold.

For unmodified assertions, this constraint is trivial. The reason is that if a given assertion meets the contextual threshold, so will all of its entailments, given that these are weaker and thus more likely to be true.[28] But with modified assertions things change; Entailment Closure will do some real work here.

Upward Monotonicity can now be derived from the fact that slifted clauses do double duty: they are asserted (albeit in a dependent sense) and they fill the argument position of the parenthetical component (which itself is independently asserted). I assume that for this reason the parenthetical component needs to be closed under the entailments of the slifted clause. We can capture this requirement by the following felicity constraint, which is a special case of (49).

(50)   ENTAILMENT CLOSURE (for slifting sentences)
       A slifting sentence E, C can be felicitously uttered only if, for any entailment C' of C, $E(C')$ meets the contextual threshold θ whenever $E(C)$ does.

This constraint derives Upward Monotonicity as follows. The main thing to notice is that the two components C and $E(C)$ associated with a slifting sentence C, E are proportional with respect to their informativity. If E is upward entailing, the strength of $E(C)$ is directly proportional to the strength of C. Thus, any weakening of C will make $E(C)$ weaker as well, and truth is always

---

[28]This is because of the fact that if p entails q, then $\mu(p) \leq \mu(q)$, and thus if $\mu(p)$ meets the threshold, so will $\mu(q)$. The notion of entailment used here is that of logical entailment. A weaker, probabilistic notion of support will not guarantee this fact. It is possible that p makes q more likely but that p has a higher unconditional probability than q.

preserved. If, however, E is downward entailing, the strength of $E(C)$ is <u>inversely</u> proportional to the strength of C. In that case, a weaker C will make $E(C)$ stronger, and truth need not be preserved.[29] In other words, only upward-entailing parentheticals preserve the entailments of the slifted clause and only such parentheticals are predicted to be felicitous.

As an illustration, imagine that (48b) above has been uttered. Since the slifted clause is asserted, speakers should be allowed to draw entailments from it, including the entailment that Snowden is a spy. However, this is ruled out because the parenthetical assertion does not guarantee that Mary doubts Snowden is a spy; it only guarantees that Mary doubts Snowden is a <u>Russian</u> spy. Since this sentence violates Entailment Closure, it is infelicitous. No such problem arises with (48a). In that case, the truth of the parenthetical assertion is preserved for all entailments of the slifted clause. In particular, if Mary thinks that Snowden is a Russian spy, it follows that she also thinks Snowden is a spy. Since Entailment Closure is upheld, the sentence is correctly predicted to be acceptable.

3.3   A broader pattern

I have proposed a general constraint on assertion that derives the upward monotonicity of slifted parentheticals. Since this constraint is expected to apply across the board, one might wonder whether its effects are only manifest in slifting sentences (which would make it look stipulative) or whether perhaps this is just the tip of the iceberg. In this subsection, I suggest that the latter is the case and list three more cases where the positive polarity status of main clauses pulls its weight.

The first case comes from a different kind of parenthetical expression. Ross (1973), Szabolcsi and Zwarts (1993), and Potts (2002) observe that an upward monotonicity restriction is operational in *as*-parentheticals as well.

(51)   (Szabolcsi and Zwarts 1993: 246)
    a.   John is our hero, as you know.
    b.   #John is our hero, as no one knows.

In addition, verb-second embedding in German is restricted in the same way. In this language, main clauses are verb second and embedded clauses are (typically) verb final. However, German also draws a contrast between regular *dass*-complements, which exhibit the usual verb-final syntax, and verb-second complements or seemingly embedded clauses with a main-clause syntax. Importantly, Scheffler (2009) observes that while attitude verbs of any monotonicity can occur in regular embedding constructions (52), verb-second embedding is restricted to upward-entailing attitude verbs (53).

(52)   a.   Peter glaubt,   dass Maria schwanger ist.
          Peter believes that Maria pregnant   is
          'Peter believes that Maria is pregnant.'
    b.   Maria bereut, dass sie nach Berlin gezogen ist.
          Maria regrets that she to   Berlin moved   is
          'Maria regrets that she moved to Berlin.'

(53)   (Scheffler 2009: 184)

---

[29]If E is non-monotone, there will be no logical relation between C and $E(C)$ and truth will not be preserved either.

a. Peter glaubt, Maria ist schwanger.
   Peter believes Maria is pregnant

   'Peter believes that Maria is pregnant.'

b. #Maria bereut, sie ist nach Berlin gezogen.
    Maria regrets she is to   Berlin moved

Finally, manner-of-speech adverbs like *truthfully* or *falsely* exhibit a similar behavior. Jackendoff (1972), Bellert (1977), and Morzycki (2016: 5.5.4) notice that positive forms can have speech act uses whereas negative forms lack such uses.

(54)   (Jackendoff 1972: 99)

a. $\left\{ \begin{array}{l} \text{Truthfully} \\ \text{Honestly} \\ \text{Sincerely} \end{array} \right\}$, I can't tell you the answer.

b. $\left\{ \begin{array}{l} \text{\#Falsely} \\ \text{\#Dishonestly} \\ \text{\#Insincerely} \end{array} \right\}$, I can't tell you the answer.

I will not attempt an analysis of these contrasts in polarity. However, it seems natural to hypothesize that they stem from a common source. If in all of these cases what is being modified is the root sentence, the monotonicity restrictions on a given modifier are likely based on a mechanism similar to the one proposed for slifting sentences.[30] If so, the empirical effects of Entailment Closure are not limited to slifting parentheticals. They form a broader pattern that cuts across empirical domains.

4   Conclusion

This paper developed a discourse model according to which an assertion updates the probabilities of common ground propositions at the level of certainty at which it is performed. Within this model, we analyzed slifting sentences as constituting complex speech acts that incorporate two linked assertions. Specifically, the parenthetical assertion was argued to provide evidence and thus modulate the strength of the assertion associated with the slifted clause.

Importantly, the proposed analysis did not reduce slifting sentences to embedding constructions and core empirical contrasts were preserved. We were able to explain why slifted claims are perceived as stronger than embedded claims by noticing that the former but not the latter have to contribute discourse-relevant information. We also derived the fact that unlike matrix clauses, slifting parentheticals are invariably upward monotone. This property was derived from the general requirement that speakers be allowed to draw entailments from what has been asserted.

The proposed analysis also taught us a good deal about the typology of assertion itself. We started off with the baseline view that asserted propositions are associated with high objective certainty, common ground status, and relevance to the question under discussion. But we soon discovered that certain forms of assertion may lack some of these canonical properties. Eventually, we arrived at a typology where regular assertions (those associated with non-parenthetical sentences)

---

[30]See Heycock (2006) for a survey of "embedded" root phenomena.

have all three ingredients, slifted assertions may fall short of the first (strength) requirement, and parenthetical assertions may lack the third (relevance) property. It was hypothesized that the lack of any of the canonical properties has to be explicitly marked in the linguistic form.

Acknowledgments

Appendix A

Proof of the Jeffrey conditionalization rule (16):

$$
\begin{aligned}
\mu(p) &= \mu(p \cap q) + \mu(p \cap \neg q) && \text{by the total probability law} \\
&= \mu(q)\,\mu_q(p) + \mu(\neg q)\,\mu_{\neg q}(p) && \text{by simple conditionalization} \\
&= \mu(q)\,\mu_q(p) + [1 - \mu(q)]\,\mu_{\neg q}(p) && \text{by the complement rule} \\
&= \theta\,\mu_q(p) + [1 - \theta]\,\mu_{\neg q}(p) && \text{by assuming } \theta = \mu(q) \\
&&& \text{QED}
\end{aligned}
$$

Proof of the three facts listed in (17):

a.
$$
\begin{aligned}
\mu_{q,1}(p) &= 1\,\mu_q(p) + [1 - 1]\,\mu_{\neg q}(p) && \text{by Jeffrey conditionalization} \\
&= \mu_q(p) && \text{by simple math} \\
&&& \text{QED}
\end{aligned}
$$

b.
$$
\begin{aligned}
\mu_{q,\mu(q)}(p) &= \mu(q)\,\mu_q(p) + [1 - \mu(q)]\,\mu_{\neg q}(p) && \text{by Jeffrey conditionalization} \\
&= \mu(q)\,\mu_q(p) + \mu(\neg q)\,\mu_{\neg q}(p) && \text{by the complement rule} \\
&= \mu(p \cap q) + \mu(p \cap \neg q) && \text{by simple conditionalization} \\
&= \mu(p) && \text{by the total probability law} \\
&&& \text{QED}
\end{aligned}
$$

c.
$$
\begin{aligned}
\mu_{q,\theta}(q) &= \theta\,\mu_q(q) + [1 - \theta]\,\mu_{\neg q}(q) && \text{by Jeffrey conditionalization} \\
&= \theta \times 1 + [1 - \theta] \times 0 && \text{by simple conditionalization} \\
&= \theta && \text{by simple math} \\
&&& \text{QED}
\end{aligned}
$$

Appendix B

Experimental and filler items for the acceptability task presented in subsection 2.3:

Lucy bought a new car, my brother said, but she actually didn't.   My brother said that Lucy bought a new car, but she actually didn't.   Lucy bought a new car, my brother said, but she might not have.   My brother said that Lucy bought a new car, but she might not have.
Tom is a very good cook, Janet thinks, but he isn't.   Janet thinks that Tom is a very good cook, but he isn't.   Tom is a very good cook, Janet thinks, but he might not be.   Janet thinks that Tom is a very good cook, but he might not be.
Peter moved to Boston, Emmy said, but I know for a fact he still lives in L.A.   Emmy said that Peter moved to Boston, but I know for a fact he still lives in L.A.   Peter moved to Boston, Emmy said, but he might still live in L.A.   Emmy said that Peter moved to Boston, but he might still live in L.A.
Richard is very smart, Mary said, but he isn't that smart really.   Mary said that Richard is very smart, but he isn't that smart really.   Richard is very smart, Mary said, but he might not be all that smart.   Mary said that Richard is very smart, but he might not be all that smart.
Steve is a film director, Nancy told us, but in reality he is just an assistant.   Nancy told us that Steve is a film director, but in reality he is just an assistant.   Steve is a film director, Nancy told us, but she might be wrong.   Nancy told us that Steve is a film director, but she might be wrong.
Steve had a career in aviation, Margaret believes, but she is wrong.   Margaret believes that Steve had a career in aviation, but she is wrong.   Steve had a career in aviation, Margaret believes, but he probably didn't.   Margaret believes that Steve had a career in aviation, but he probably didn't.
Paul is Mexican, my friend believes, but in fact he is from Puerto Pico.   My friend believes that Paul is Mexican, but in fact he is from Puerto Rico.   Paul is Mexican, my friend believes, but he might be from Puerto Rico.   My friend believes that Paul is Mexican, but he might be from Puerto Rico.
Melissa trained every day, her coach claimed, but she actually didn't.   Her coach claimed that Melissa trained every day, but she actually didn't.   Melissa trained every day, her coach claimed, but I doubt she did.   Her coach claimed that Melissa trained every day, but I doubt she did.

It's done, and if it's done it's done.
Kurt lives in Paris but he doesn't like it there.
The sun rises in the east and sets in the west.
Birds can fly, although some of them cannot.
Adam is Canadian but he is also Russian.
The book reads well but no one has ever read it.
Matt went back to Australia but he has never been there before.
It's raining but it's not raining.

References

Aikhenvald, A. (2004). *Evidentiality*. Oxford: Oxford University Press.

AnderBois, S. (2016). Semantics and pragmatics of (not-)at-issueness in Yucatec Maya attitude reports. *Semantics & Pragmatics 9*(19), 1–55.

Asher, N. (1987). A typology for attitude verbs and their anaphoric properties. *Linguistics and Philosophy 10*(2), 125–197.

Asher, N. (2000). Truth conditional discourse semantics for parentheticals. *Journal of Semantics 17*, 31–50.

Asher, N. and A. Lascarides (2003). *Logics of Conversation*. Cambridge, UK: Cambridge University Press.

Bach, K. (1999). The myth of conventional implicature. *Linguistics and Philosophy 22*, 327–366.

Bach, K. and R. M. Harnish (1979). *Linguistic Communication and Speech Acts*. Cambridge, Massachusetts/London, England: MIT Press.

Beaver, D. I., C. Roberts, M. Simons, and J. Tonhauser (2017). Questions Under Discussion: Where information structure meets projective content. *Annual Review of Linguistics 3*, 265–284.

Bellert, I. (1977). On semantic and distributional properties of sentential adverbs. *Linguistic Inquiry 8*(2), 337–351.

Bradley, S. (2019). Imprecise probabilities. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. URL = https://plato.stanford.edu/archives/spr2019/entries/imprecise-probabilities/.

Brown, J. and H. Cappelen (2004). Assertion: an introduction and overview. In J. Brown and H. Cappelen (Eds.), *Assertion: New Philosophical Essays*, pp. 1–17. Oxford: Oxford University Press.

Büring, D. (2003). On d-trees, beans, and b-accents. *Linguistics and Philosophy 26*(2), 511–545.

Culicover, P. W. (1992). English tag questions in Universal Grammar. *Lingua 88*, 193–226.

Davis, C., C. Potts, and M. Speas (2007). The pragmatic values of evidential sentences. In *Proceedings of Semantics and Linguistic Theory 17*, pp. 71–88.

DeRose, K. (2004). Assertion, knowledge, and context. *The Philosophical Review 111*(2), 167–203.

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

Grimshaw, J. (2011). The place of slifting in the English complement system. In *Generative Initiatives in Syntactic Theory 3*, Handout, pp. 246–255.

Groenendijk, J. and M. Stokhof (1984). On the semantics of questions and the pragmatics of answers. In F. Landman and F. Veltman (Eds.), *Varieties of Formal Semantics*, pp. 143–170. Dordrecht: Foris.

Haddican, B., A. Holmberg, H. Tanaka, and G. Tsoulas (2014). Interrogative slifting in English. *Lingua 138*, 86–106.

Halpern, J. Y. (2003). *Reasoning about Uncertainty*. Cambridge, Massachusetts/London, England: MIT Press.

Heycock, C. (2006). Embedded root phenomena. In H. van Riemsdijk and M. Everaert (Eds.), *The Blackwell Companion to Syntax, vol. II*, pp. 174–209. Oxford: Blackwell.

Hindriks, F. (2007). The status of the knowledge account of assertion. *Linguistics and Philosophy 30*, 393–406.

Hooper, J. B. (1975). On assertive predicates. In J. Kimball (Ed.), *Syntax and Semantics 4*, pp. 91–124. Academic Press.

Huddleston, R. and G. Pullum (2002). *The Cambridge Grammar of the English Language*. Cambridge, UK: Cambridge University Press.

Hunter, J. (2016). Reports in discourse. *Dialogue & Discourse 7*(4), 1–35.

Jackendoff, R. (1972). *Semantic Interpretation in Generative Grammar*. Cambridge, Massachusetts/London, England: MIT Press.

Jayez, J. and C. Rossari (2004). Parentheticals as conventional implicatures. In F. Corblin and H. de Swart (Eds.), *Handbook of French Semantics*, pp. 211–229. Stanford: CSLI.

Jeffrey, R. (1990). *The Logic of Decision*. Chicago: The University of Chicago Press.

Koev, T. (2018). Notions of at-issueness. *Language and Linguistics Compass 12*(12), 1–16.

Lackey, J. (2007). Norms of assertion. *Mind 41*, 594–626.

Ladusaw, W. A. (1980). *Polarity Sensitivity as Inherent Scope Relations*. Unpublished dissertation, University of Texas, Austin.

Lewis, D. (1976). Probabilities of conditionals and conditional probabilities. *The Philosophical Review 85*, 297–315.

MacFarlane, J. (2005). Making sense of relative truth. In *Proceedings of the Aristotelian Society 105*, pp. 321–339.

MacFarlane, J. (2011). What is assertion? In J. Brown and H. Cappelen (Eds.), *Assertion: New Philosophical Essays*, pp. 79–96. Oxford: Oxford University Press.

Maier, E. and C. Bary (2015). Three puzzles about negation in non-canonical speech reports. In *Proceedings of the 20th Amsterdam Colloquium*, pp. 246–255.

McCready, E. (2015). *Reliability in Pragmatics*. Oxford: Oxford University Press.

McCready, E. and N. Ogata (2007). Evidentiality, modality and probability. *Linguistics and Philosophy 30*, 147–206.

Morzycki, M. (2016). *Modification*. Cambridge, UK: Cambridge University Press.

Moss, S. (2015). On the semantics and pragmatics of epistemic vocabulary. *Semantics & Pragmatics 8*(5), 1–81.

Murray, S. (2014). Varieties of update. *Semantics & Pragmatics 7*(2), 1–53.

Potts, C. (2002). The syntax and semantics of as-parentheticals. *Natural Language & Linguistic Theory 20*, 623–689.

Potts, C. (2005). *The Logic of Conventional Implicatures*. Oxford: Oxford University Press.

Rooryck, J. (2001). Evidentiality I. *Glot International 5*(4), 125–133.

Ross, J. R. (1973). Slifting. In M. H. M. Gross and M. P. Schützenberger (Eds.), *The Formal Analysis of Natural Language*, pp. 133–169. The Hague: Mouton.

Rudin, D. (2018). Uncertainty and persistence: A Bayesian update semantics for probabilistic expressions. *Journal of Philosophical Logic 47*(5), 365–405.

Scheffler, T. (2009). Evidentiality and German attitude verbs. In *Proceedings of Penn Linguistics Colloquium 32*, pp. 183–192.

Searle, J. R. (1969). *Speech Acts: An Essay in the Philosophy of Language*. Cambridge: Cambridge University Press.

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

Simons, M., J. Tonhauser, D. Beaver, and C. Roberts (2010). What projects and why. In *Proceedings of Semantics and Linguistic Theory 20*, pp. 309–327.

Slote, M. (1975). Assertion and belief. In J. Dancy (Ed.), *Papers on Language and Logic: Proceedings of the Conference on the Philosophy of Language and Logic*, pp. 94–102. Keele University Library.

Stalnaker, R. (1978). Assertion. In P. Cole (Ed.), *Syntax and Semantics, Volume 9: Pragmatics*, pp. 315–332. New York: Academic Press.

Stenius, E. (1967). Mood and language-game. *Synthese 17*, 254–274.

Szabolcsi, A. and F. Zwarts (1993). Weak islands and an algebraic semantics for scope taking. *Natural Language Semantics 1*, 235–284.

Unger, P. (1975). *Ignorance: A Case for Scepticism*. Oxford: Oxford University Press.

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

Weiner, M. (2007). Norms of assertion. *Philosophy Compass 2*(2), 187–195.

Weiner, M. (2013). Must we know what we say? *The Philosophical Review 114*(2), 227–251.

Williamson, T. (2004). Knowing and asserting. *The Philosophical Review 105*(4), 489–523.

Yalcin, S. (2010). Probability operators. *Philosophy Compass 5*(10), 916–937.