

Strong beliefs, weak commitments¹

Todor KOEV — *University of Konstanz*

Abstract. The standard Hintikkan semantics views *believe* as a universal quantifier over possible worlds (Hintikka, 1969). This semantics (i) fails to capture the fact that *believe* is gradable (cf. *partially believe* or *fully believe*) and (ii) makes no predictions about the degree of certainty of the belief agent toward the prejacent. To remedy these problems, I propose a scalar semantics along the lines of Kennedy and McNally’s (2005) analysis of gradable adjectives, arguing that *believe* is a maximum-degree predicate. While belief attributions are sometimes interpreted as hedges (e.g., *I believe it’s raining* can be taken as a statement of uncertainty), I point out that such uses are restricted to contexts in which the belief component is not relevant to the question under discussion. Following up on a suggestion made in Chemla (2008), I propose that the weak sense of *believe* arises as an antipresupposition, a scalar inference derived through competition with a presuppositionally stronger *know*-competitor. Contra Hawthorne et al. (2016), I argue that the intuition of weakness is due not to reduced modal force but rather to the subjectivity of modal content, amounting to a situation in which the agent has full subjective confidence in the prejacent but fails to publicly commit to it.

Keywords: belief, modality, gradability, subjectivity v. objectivity, antipresupposition, questions under discussion.

1. The Hintikkan orthodoxy and its problems

The verb *believe* plays a pivotal role in semantic research as it underlies a number of widely studied phenomena, e.g., opacity, presupposition projection, neg-raising, the norms of assertion, etc. It is then crucial to understand its core interpretational properties. In formal semantics, it has become standard to analyze *believe* as a universal quantifier over possible worlds. Ever since Hintikka (1969), a belief attribution is taken to state that the **prejacent** (=the embedded proposition) is true in all of the agent’s doxastic alternatives. This is usually rendered as follows, where $Dox_{x,w}$ stands for the set of x ’s doxastic alternatives in w , i.e., the set of possible worlds compatible with everything x believes in w .

$$(1) \quad \llbracket \text{believe} \rrbracket^w = \lambda p \lambda x. \forall w' \in Dox_{x,w} : p(w')$$

Though very popular, this Hintikkan orthodoxy fails on at least two counts.

- **Gradability:** *Believe* is a gradable predicate, as evidenced by its ability to tolerate degree modification (cf. *partially believe* or *fully believe*). But it remains unclear how this gradability property can be modeled if the force of *believe* is fixed by a quantifier once and for all. Ideally, the strength of *believe* should be able to be manipulated the same way degree morphology manipulates the degree argument of gradable adjectives. Yet a simple quantificational semantics does not supply a degree argument.

¹I would like to thank Danny Fox, Louise McNally, Maša Močnik, Roger Schwarzschild, and the audiences at *SuB*, the University of Cologne, the University of Konstanz, and MIT. For English judgments, I am indebted to Kurt Erbach, James Griffiths, Kyle Rawlins, and Peter Sutton.

- Commitment strength: The standard semantics makes no predictions as to how strongly unmodified uses of *believe* commit its agent to the prejacent. Does *Fred believes it's raining* require that Fred is certain it is raining or just that Fred finds it likely it is raining? The lexical entry above states that all doxastic alternatives are worlds in which the prejacent is true. But universal force alone does not entail a specific degree of certainty without a clear idea of how the agent is linked to the domain of quantification. Since the set of doxastic alternatives is defined as the set of worlds compatible with everything the agent “believes” in the world of evaluation, the issue of commitment strength is but shifted to the metalanguage. It is not derived and needs to be stipulated.

Quite surprisingly, these two issues have rarely been put on the table in formal semantics, although similar questions about modal adjectives have been discussed in work on graded modality (Kratzer, 1991; Portner, 2009; von Stechow and Gillies, 2010; Yalcin, 2010; Klecha, 2014; Lassiter, 2017). The gradability issue can be addressed if the semantics of *believe* is recast in a degree-based framework as the one developed in Cresswell (1976), von Stechow (1984), Kennedy (1999), Kennedy and McNally (2005), a.o. Doing so appears to posit less of a challenge and the first steps in this direction have already been made (e.g., Lassiter, 2017; Santorio and Romoli, 2017). The commitment strength issue, however, has barely been touched upon in the literature, with the notable exception of Hawthorne et al. (2016). Deciding on the strength of *believe* turns out to be a challenging task, and the bulk of this paper will be devoted to providing a plausible answer.

Is *believe* weak or strong? The empirical picture appears to be varied right off the bat. For example, an utterance of (2) can be taken as a description of certainty or as a hedge toward the embedded proposition, depending on the context. The former (strong) reading will be brought out if we are playing a game where everyone is required to list some of their beliefs. The latter (weak) reading may come about if we are making guesses about who the next president of the U.S. is going to be.

(2) I believe Oprah will win the next election.

Given this flexibility, there are two plausible views on the strength of *believe*. The first view is that this verb implies full certainty, i.e., the agent regards the embedded proposition as true. I call this view **Strong Believers** (SB) and state it more formally as in (3). Here $Cr_x(p)$ stands for x 's credence (or subjective confidence) in p .

(3) Strong Believers
 $\llbracket \text{believe} \rrbracket(p)(x)$ iff $Cr_x(p) = 1$

Though rarely articulated, SB seems to be what has been implicitly assumed all along. For example, in a rare moment of explicitness, Lasnik (2005: 675) states: “To believe something is to consider it true.” SB appears to be the prevalent view in philosophy as well (e.g., Levi, 1991; Williamson, 2002; Clark, 2013).

The second view is what I call **Weak Believers** (WB). This view states that the subjective probability associated with *believe* exceeds some contextually specified threshold, which would

typically be 50% but can be significantly influenced by the context.

- (4) Weak Believers
 $\llbracket \text{believe} \rrbracket(p)(x)$ iff $Cr_x(p) > \theta_{bel}$

WB is defended in Hawthorne et al. (2016), who offer a number of empirical arguments in its support (see also Foley, 1992; Lauer, 2017). I address these arguments in section 2, along with a range of fresh data that argue against this view.

It is also possible to claim that *believe* is lexically ambiguous between a weak and a strong interpretation, so that both WB and SB are valid. I have two worries about such a proposal, one conceptual and one empirical. The conceptual worry stems from the fact that the alleged readings are logically dependent, as the stronger entails the weaker. This way of cutting the meaning pie is hardly theoretically parsimonious and it is unclear how the language learner can acquire such a distinction. The empirical worry is that under an ambiguity approach, one would expect *I believe_{weak} he is going to win but I don't believe_{strong} he will* to be as felicitous as *It's likely he is going to win but it's not certain*. In reality, the latter sentence is acceptable while the former sentence sounds contradictory. Given these obvious difficulties, I see no merit in pursuing the ambiguity approach any further.

The main claim of this paper is that *believe* is a maximum-degree gradable modal. This means that *believe* makes available a degree argument that can be manipulated by degree morphology, yet in the absence of (overt) degree morphology it refers to the top of the scale. Crucially, *believe* differs from modal adjectives like *certain* or *sure* in that it encodes a measure of subjective rather than objective probabilities (hence the use of a credence function in (3)-(4) above). In other words, *believe* is subjective but strong: the agent regards the prejacent as true but may lack sufficient evidence for it. This opens up the possibility that a belief agent has full confidence in the prejacent but does not want to go on record and publicly endorse it, thus giving rise to the intuition of weakness. I will argue that the intuition of weakness arises as an antipresupposition, i.e., a kind of scalar inference derived through competition with a presuppositionally stronger *know*-competitor. While this view supports SB over WB, it also emphasizes the idea that the perceived weakness of *believe* is not tied to its modal force but rather to its modal content. Once we acknowledge that *believe* invokes subjective certainty, we can have an explanation for why it behaves as a strong modal throughout while at the same can serve as a hedge. The hedging use does not weaken its strong subjective force but rather hints at the lack of objective certainty.

The structure of the paper is as follows. Section 2 lays out the empirical landscape on the gradability and the strong nature of *believe*. Section 3 develops a scalar semantics for *believe*, one that derives its gradability, its closure under conjunction, and its strong force, yet leaves the door open to hedging uses. Section 4 is the conclusion.

2. The empirical picture

This section sifts through five sets of data that furnish converging evidence in favor of viewing *believe* as a gradable predicate that takes a maximum-degree standard. Most importantly, I demonstrate that the weak interpretational component becomes visible only when brought to the fore by contextual factors.

2.1. Gradability

Believe is a gradable predicate. It can participate in comparative (5) and equative constructions (6) (if mediated through gradable adverbs like *strongly*), and can be directly modified by minimality (7), maximality (8), and proportional modifiers (9). (All examples below are culled from the web.)

- (5) a. He believes more strongly than I do that the organization of the executive branch of the federal government matters a great deal.
 b. Men believe less strongly than women that they have control over their future health or that personal actions contribute to good health.
- (6) Each [farmer] believes as strongly as the other that his crops will not survive another week without water, and each cares as much as the other about the survival of his crops.
- (7) Atticus partially believes that prejudice exists because people do not understand each other [...].
- (8) a. I strongly believe that life is too short to eat mediocre meals.
 b. Darwin says he almost believes that species are not immutable.
 c. Theresa fully believes that we all have the inner ability to achieve what we desire and sometimes it takes input from others to kick start that process.
 d. I truly believe, 100 percent, that for every person reading this article, I can go one by one and determine your potential for success by looking at only two basic principles.
- (9) a. One possibility is that Charles half believes that there is a real danger, and that he is, literally, at least half afraid.
 b. This has taken me lots of research to come to this conclusion, but I believe 95 percent that it is.

The fact that *believe* is gradable does not prejudge the choice between SB and WB, i.e., the question of whether an unmodified use of *believe* implies full or partial certainty. The reason is that – when occurring outside a degree construction – gradable predicates may pick different standards of comparison. Unger (1971) was the first to distinguish between two kinds of gradable adjectives, depending on how the standard is chosen. The standard for **relative** adjectives like *tall* is selected contextually (typically taken from the middle of the scale), while the standard for **absolute** adjectives like *bent*, *certain*, *full* is fixed as the minimum or the maximum of the scale. Kennedy and McNally (2005) and Kennedy (2007) convincingly argue that the relative/absolute distinction boils down to differences in scale structure. They classify gradable predicates depending on whether the associated scale is open or closed on its ends, producing the following typology.

- | | | | | |
|------|----|-------------------------------|--------|---------------------------|
| (10) | a. | totally open scale | ○————○ | <i>tall, short</i> |
| | b. | lower-closed upper-open scale | ●————○ | <i>bent, straight</i> |
| | c. | lower-open upper-closed scale | ○————● | <i>certain, uncertain</i> |
| | d. | totally closed scale | ●————● | <i>full, empty</i> |

This scale typology is empirically supported by the distribution of degree modifiers. Thus, proportional modifiers (e.g., *half* or *mostly*) are only acceptable with adjectives encoding totally

closed scales; maximum-degree modifiers (e.g., *perfectly*) are only compatible with adjectives encoding upper-closed or totally closed scales; and minimum-degree modifiers (e.g., *slightly*) can only occur with adjectives encoding lower-closed or totally closed scales.

Kennedy and McNally draw the following important generalization: in their unmarked or “positive” form, adjectives associated with totally open scales take relative standards while adjectives associated with partially or totally closed scales take absolute standards. For example, *tall* is associated with a totally open scale (it cannot occur with minimality or maximality modifiers like *slightly* or *completely*, at least not without coercion) and takes a relative standard (what counts as tall will depend on the comparison class at hand). By contrast, *full* has an upper-closed scale (things can be said to be completely full) and takes the maximal degree of the scale as its standard (a bottle is said to be full if it cannot accept more liquid, modulo pragmatic slack).²

Given the possibility for independently probing into scalar structure, we can use the Kennedy/McNally generalization to assess whether *believe* is a relative or an absolute predicate. What kind of scale is *believe* associated with? The data in (5)-(9) clearly argue for a totally closed scale, due to the compatibility with minimum, maximum, and proportional modifiers. I will thus assume that *believe* has the ratio scale [0, 1], which it shares with probability measures. Since its scale is totally closed, the Kennedy/McNally generalization predicts that *believe* picks out an absolute rather than a relative standard. This prediction is in line with SB, which entails an absolute standard of 1 (full subjective certainty), but not compatible with WB, which does not set a fixed standard. We can conclude that *believe* takes an absolute standard, not a relative one.

2.2. Missing quantity implicatures

WB and SB make different predictions about potential quantity implicatures associated with unmodified belief sentences. If *believe* is weak, its use is expected to routinely trigger implicatures to the effect that the agent has doubts about the prejacent. By contrast, if *believe* is strong, no generation of such implicatures is predicted. The lack of systematic quantity implicatures associated with belief attributions in either root or embedded positions is a first hint that the latter view is more on the right track.³

- (11) a. Kamala believes that America needs universal health care.
 b. $\not\rightarrow$ Kamala is not (fully) convinced that America needs universal health care.

²The explanation for the first part of Kennedy and McNally’s generalization is straightforward: if a scale lacks endpoints, an adjective associated with it needs contextual support in order to find an appropriate standard. The explanation for the second part of the generalization requires an additional optimization principle. Kennedy (2007) thus invokes Interpretive Economy, according to which truth conditions favor conventional meaning over context. Given this principle, if a scale provides endpoints, an adjective must use these when picking a standard before it looks for it elsewhere, i.e., before it involves the context.

³In these examples the putative implicature associated with *x believes p* is rendered as *x is not convinced that p* rather than, say, as *x is not certain that p*. The reason is that, according to the current proposal, *believe* encodes a subjective measure, and *convinced* seems to share this property with it. While *x believes p but is not certain that p* is fully natural, I argue that *certain* (unlike *believe*) is associated with an objective measure and the contrast established in sentences of this shape is about the degree of public commitment rather than modal force (see subsection 3.3 for details).

- (12) a. According to the press, Kamala believes that America needs universal health care.
 b. ✗ According to the press, Kamala is not (fully) convinced that America needs universal health care.

Attributing belief to attitude holders whose subjective certainty level is fixed by the context also recommends SB over WB. Under a weak story, the utterance in (13) should be ruled out as underinformative, and yet it is judged to be felicitous. Conversely, (14) demonstrates that it is hard to sincerely attribute belief to an agent who doubts the prejacent to some degree. WB lets us expect that the suggested inference can be drawn, given that the attitude holder finds the prejacent likely to be true. By contrast, SB blocks this inference due to lack of full (subjective) certainty.

- (13) *Context: We don't know whether the transfer student passed the midterm but Jill has no doubts he did.*
 Jill believes the transfer student passed the midterm.
- (14) Mueller finds it likely but still has some doubts that Russia influenced the election.
 #So, Mueller believes that the Russians did it.

The last kind of data I discuss in this subsection involves scalar gradation. If *believe* is strong, it should be able to strengthen a weaker modal with the same content but it should not allow for itself to be further strengthened by a stronger modal with the same content. Starting with the former prediction, Hawthorne et al. (2016: 1398) observe that *believe* cannot strengthen *think*, which they view as weak.

- (15) ??Tim thinks it's raining, but he doesn't believe that it is.

However, I see no reason not to adopt the view that *think* is as strong as *believe*. After all, the two verbs are nearly synonymous and have a similar distribution, with the important caveat that *think* is not gradable (cf. **partially think* or **fully think*). If this is on the right track, *think* is strong but subjective, so (15) is expected to be out. (A lexical entry for *think* that has this profile is given in subsection 3.1.)

Perhaps the most well-known objection to the second prediction is the old dictum that knowledge entails belief (but not vice versa), supported by natural gradations as in the following example.

- (16) Scientists believe there is water on Mars. In fact, they know it.

However, it has also been argued that while *know* is stronger than *believe* in terms of its presuppositions and perhaps in terms of content as well, it is not stronger in force (see Chemla, 2008; Sauerland, 2008; Schlenker, 2012). *Know* is presuppositionally stronger in that it adds the factive inference that the prejacent is true and it may also be truth-conditionally stronger in that it entails that the attitude holder has appropriate evidence for the prejacent. Either property could be invoked to explain the felicitous gradation above without ascribing a raised subjective certainty to the agent. To put it differently, the modal gradation above need not raise the likelihood of the prejacent from the point of view of the attitude agents; it may do so merely from the point of view of the speaker.

2.3. Closure under conjunction

One difference between strong and non-strong modals concerns the way they interact with conjunction. Strong modals are closed under conjunction, i.e., they license the entailment in (17), where M ranges over modals and p, q range over propositions. This is illustrated in (18). Non-strong modals, on the other hand, do not have this property, see (19)-(20).

$$(17) \quad M(p) \wedge M(q) \models M(p \wedge q)$$

(18) a. It's certain that Sean is in Rome and it's certain that he is catholic.

b. \models It's certain that Sean is in Rome and that he is catholic.

(19) *Each week Jack spends (in no particular order) 3 nights at the local pub and gets drunk, 2 nights at the same pub but stays sober, and 2 nights at home where he also gets drunk. On a given night, I say:*

a. Jack is probably at the pub. True (chance = 5/7)

b. Jack is probably drunk. True (chance = 5/7)

c. Jack is probably at the pub drunk. False (chance = 3/7)

(20) a. It's possible Jane is in Italy and it's also possible Jane is in France.

b. $\not\models$ It's possible Jane is in Italy and in France.

Since (21) presents a valid entailment, *believe* pairs up with strong modals in this respect. Notice that even a hedged reading of *believe* licenses this type of inference (22). This fact suggests that the intuition of weakness does not translate into non-strong modal force.

(21) a. Ron believes Mia is hawt and he also believes she is going to marry him.

b. \models Ron believes that Mia is hawt and that she is going to marry him.

(22) a. I believe Sean is catholic, but I'm not sure.

b. I believe Sean is in Rome, but I'm not sure.

c. \models I believe Sean is catholic and he is Rome, but I'm not sure.

There is a lot of discussion in the philosophical literature about whether the beliefs of a rational agent are closed under conjunction. While most philosophers agree this should be so (e.g., Hintikka, 1962; Levi, 1973; Leitgeb, 2014), detractors point out that the closure property leads to the **lottery paradox** (Kyburg, 1961; Foley, 1992). A classical version of the lottery paradox for rational belief goes as follows. Consider a fair lottery with 100 tickets and one winner. It seems rational to believe the statement "Ticket #1 will not win", as it has a solid 99% chance of being true. But the same goes for the statements "Ticket #2 will not win", "Ticket #3 will not win", and so on down the line up to "Ticket #100 will not win". By the closure property, it should then be rational to believe the statement "No ticket will win". But this contradicts the assumption that one ticket will win.

The important thing to notice here is that the lottery paradox is about the norms of rational belief rather than the semantics of the verb *believe* per se. If rational belief is understood as reaching some high but non-maximal level of confidence, it is indeed reasonable to reject the closure property. However, when the paradox is brought to bear on the use of the verb *believe*, the linguistic judgments can be disputed, as Hawthorne et al. (2016) themselves admit. In the lottery scenario, it may not be entirely sincere to assert that one believes that one's ticket will

lose based on numeric probabilities alone. Indeed, if we assume that *believe* is strong, any shortage of (subjective) certainty would make the premises false and the paradox would not arise.

2.4. Neg-raising and modal strength

Neg-raising is a phenomenon whereby a matrix negation is interpreted as if it takes scope inside an embedded clause, so that *x doesn't believe p* comes to mean *x believes not p* (Bartsch, 1973; Horn, 1989; Gajewski, 2007; Romoli, 2013; Homer, 2015). While there is no universally accepted analysis of neg-raising, semantic accounts typically cash in on Bartsch's "excluded middle" principle, according to which the agent holds the described attitude towards the embedded proposition or its negation. Thus, if *x doesn't believe p* is uttered and *x believes p* \vee *x believes not p* is assumed to hold, we can conclude *x believes not p*, since the assertion is only compatible with the second disjunct of the excluded middle principle. The strengthened neg-raised reading of the original utterance is now derived.

Since *believe* is a neg-raising predicate, a legitimate question to ask is what other predicates fall in the same class and whether there are generalizations to be drawn about its members. Hawthorne et al. (2016) hypothesize that neg-raising occurs with weak modal predicates (e.g., *think*, *want*, *like*, *advise*) but not with strong modal predicates (e.g., *know*, *need*, *love*, *order*). The fact that *believe* shares this property with the former group, they argue, bespeaks a weak semantics.

However, a closer look reveals that the alleged link between weak modal force and neg-raising is not supported by the data. Horn (1989: ch.5) draws a distinction between three types of modals depending on their perceived strength: weak scalars (e.g., *possible*, *allowed*), mid-scalars (e.g., *likely*), and strong scalars (e.g., *certain*, *necessary*). Taking into consideration the crosslinguistic picture, he goes on to show that weak scalars never license neg-raising, mid-scalars typically do, and strong scalars may or may not license it. Thus, with the exception of weak scalars, modal strength is not a reliable indicator of neg-raising behavior. Notice, for example, that *want* is a neg-raising predicate while *desire* is not (Gajewski, 2007), but it is not obvious at all that the two verbs differ in strength. Thus, the fact that *believe* licenses neg-raising does not decide on its modal strength, except for excluding the possibility that it is a weak scalar.

2.5. Hedging

The main challenge to SB stems from the observation that a statement of belief can be used as a hedging device. The belief agent below explicitly disavows responsibility for the truth of the prejacent.⁴

(23) I believe it's raining, but I'm not sure it's raining. (Hawthorne et al., 2016)

⁴As Hawthorne et al. (2016) point out, some English speakers have a preference for *think* over *believe* in these examples, which could significantly weaken their argument. However, my impression is that most speakers do accept hedging examples like these, at least as root sentences. In embedded positions, such structures may be less acceptable (cf. *?Suppose I believe it's raining but I'm not sure*).

I will argue below that such examples do not put into question the strong force of *believe* because all they do is establish a contrast between privately held convictions and publicly expressed commitments. Under this view, (23) means somethings like “The speaker is fully confident that it is raining but she does not want to publicly commit to it (presumably because she lacks sufficient evidence)”. Once we acknowledge that doxastics like *believe* or *think* differ from modal adjectives like *certain* or *sure* in that the former invoke subjective certainty, we can understand why a strong subjective modal can be used as a hedge on the objective certainty of the agent.

It is important to point out that the hedging use comes with relevance restrictions on it. Hedging uses turn out to be sensitive to what the conversation is about, or the **question under discussion** (Ginzburg, 1996, 2012; Roberts, 1996; Büring, 2003; Farkas and Bruce, 2010). As the data below shows, the intuition of weakness arises when the prejacent (rather than the belief component) is at issue (24). When what is at issue is the belief component itself, the weak component of *believe* is difficult to access (25).

- (24) {Is capitalism better than socialism?}
I believe so (but I’m not sure).
- (25) {Tell us about your political beliefs.}
I believe capitalism is better than socialism (?but I’m not sure).

In summary, we see that in order for *believe* to appear weak, the belief attribution itself must not be relevant in the current discourse. This observation suggests that the intuition of weakness should not be baked into the lexical entry of this verb. The weak component is more likely a by-product of the way this verb interacts with contextual factors.

3. A strong, scalar, and subjective semantics for *believe*

Given the discussion in the previous section, we need a semantics that derives the following facts.

- *Believe* is gradable; its strength can be manipulated by degree modifiers like *partially*, *truly*, or *95 percent*. In its unmodified use, it is an absolute predicate that takes a maximum-degree standard.
- *Believe* is closed under conjunction with respect to its internal argument: *x believes p* and *x believes q* jointly entail *x believes (p and q)*.
- Being maximum degree, *believe* is strong: an unmodified use entails full certainty on the part on the attitude holder. At the same time, the implied commitments can be weak, provided that the belief component is not relevant to the question under discussion.

This section offers a strong semantics for *believe* that allows for degree modification, entails the closure property, and derives the intuition of weakness as a particular kind of a scalar inference.

3.1. *Believe* as maximum degree

In order to capture the gradability of *believe*, I follow the approach to gradable adjectives proposed in Cresswell (1976), von Stechow (1984), Kennedy (1999), Kennedy and McNally (2005), and extended to modal adjectives in Lassiter (2017) and Santorio and Romoli (2017). The entry below states that the belief agent's credence in the prejacent meets some threshold, where Cr is a credence function and p, d, x are variables over propositions, degrees, and individuals (respectively).⁵

$$(26) \quad \llbracket \text{believe} \rrbracket = \lambda p \lambda d \lambda x. Cr_x(p) \geq d$$

There are several important things to notice about this meaning. The first is that *believe* expresses credences, i.e., subjective probabilities. This will be important in subsection 3.3 in order to ensure that *believe* can tolerate weak public commitments. The proposed semantics also entails that *believe* is associated with the probabilistic scale $[0, 1]$, which is fully closed (see subsection 2.1 for empirical arguments). Given Kennedy's (2007) principle of Interpretive Economy (mentioned in footnote 2 above), we predict that *believe* is an absolute predicate. Finally, notice that the degree argument is filled after the propositional argument. While degree arguments may be assumed to be fed in different orders, setting things up this way will allow us to maintain single lexical entries for degree modifiers that work across the adjectival and the verbal domain.

The degree argument of gradable adjectives is always filled by degree morphology. For positive forms, where no overt degree morpheme is present, it has been assumed that the norm of comparison is supplied by a covert morpheme called *pos*. Cresswell (1976), who pioneered this idea, assumed that *pos* contributes something like "more than average", so that *Bill is tall* comes to mean that Bill is of an above average height. I thus assume the following semantics for *pos* (adopted from Kennedy and McNally, 2005), where C is a contextually supplied comparison class of appropriate objects.

$$(27) \quad \llbracket \text{pos} \rrbracket^C = \lambda P \lambda x. \exists d [\mathbf{standard}(d, P, C) \wedge P(d)(x)]$$

Kennedy and McNally (2005) suggest that the first conjunct in the above formula links the degree argument to the right type of standard depending on the features of the selected predicate and relative to a comparison class. This conjunct is spelled out as follows, assuming that μ_P is the measure function associated with P and s_P is the scale associated with P .

$$(28) \quad \mathbf{standard}(d, P, C) = \begin{cases} d > \mathbf{avg}\{\mu_P(x) \mid x \in C\} & \text{if } P \text{ is relative} \\ d > \mathbf{min}(s_P) & \text{if } P \text{ is minimum degree} \\ d = \mathbf{max}(s_P) & \text{if } P \text{ is maximum degree} \end{cases}$$

The standard norm for relative adjectives is context sensitive; it is the average degree of P -ness of the objects inside some contextually specified comparison class. By contrast, the standard norms for absolute adjectives do not depend on the facts in the world. They are fixed as the minimum or the maximum of the relevant scale.

Let us assume that *pos* attaches to unmodified VPs headed by *believe*, as below. Since in this

⁵Following up on the discussion in subsection 2.2, I suggest that *think* has a similar semantics but lacks a degree argument: $\llbracket \text{think} \rrbracket = \lambda p \lambda x. Cr_x(p) = 1$.

case *pos* equates the degree of belief to the maximum of the probability scale (=1), we derive the fact that unmodified *believe* is strong.

- (29) *pos* [believes it's raining]
 a. $\llbracket \text{believes it's raining} \rrbracket = \lambda d \lambda x. Cr_x(\llbracket \text{rain} \rrbracket) \geq d$
 b. $\llbracket \text{pos [believes it's raining]} \rrbracket = \lambda x. \exists d [d = 1 \wedge Cr_x(\llbracket \text{rain} \rrbracket) \geq d]$
 $= \lambda x. Cr_x(\llbracket \text{rain} \rrbracket) = 1$

We can adopt the same kind of meanings for overt degree modifiers that usually combine with gradable adjectives. The following lexical entries are modeled on Kennedy and McNally (2005) and preserve the logical types. For presentation purposes, I take one instance of minimality, maximality, and proportional modifiers each.

- (30) a. $\llbracket \text{partially} \rrbracket = \lambda P \lambda x. \exists d [d > \mathbf{min}(s_P) \wedge P(d)(x)]$
 b. $\llbracket \text{truly} \rrbracket = \lambda P \lambda x. \exists d [d = \mathbf{max}(s_P) \wedge P(d)(x)]$
 c. $\llbracket \text{95 percent} \rrbracket = \lambda P \lambda x. \exists d [d = 0.95 \wedge P(d)(x)]$

For example, belief sentences with minimality modifiers like *partially* give rise to the meaning composition shown below.

- (31) partially [believes eating pizza is healthy]
 a. $\llbracket \text{believes eating pizza is healthy} \rrbracket = \lambda d \lambda x. Cr_x(\llbracket \text{eat pizza healthy} \rrbracket) \geq d$
 b. $\llbracket \text{partially [believes eating pizza is healthy]} \rrbracket$
 $= \lambda x. \exists d [d > 0 \wedge Cr_x(\llbracket \text{eat pizza healthy} \rrbracket) \geq d]$
 $= \lambda x. Cr_x(\llbracket \text{eat pizza healthy} \rrbracket) > 0$

Comparative morphology in belief sentences is more challenging. The reason is that *believe* does not directly compose with comparative morphemes the way gradable adjectives or verbs like *love* or *hate* do. Instead, beliefs are juxtaposed by comparative forms of gradable adverbs like *strongly*. This necessitates additional entries for degree modifiers that compose with adverbs rather than adjectives. To show how this can be done, let us first posit the following semantically bleached meaning for *strongly*.

- (32) $\llbracket \text{strongly} \rrbracket = \lambda P \lambda d \lambda x. P(d)(x)$

The “strong” meaning of unmodified *strongly* comes not from the adverb itself but from the fact that it composes with an adverbial version of *pos*. This is demonstrated below, where d_s is the relevant threshold for beliefs that count as strong.

- (33) $\llbracket \text{pos}_{Adv} \rrbracket^C = \lambda A \lambda P \lambda x. \exists d [\mathbf{standard}(d, P, C) \wedge A(P)(d)(x)]$
 (34) [*pos_{Adv} strongly*] [believes *S*]
 a. $\llbracket \text{pos}_{Adv} \text{ strongly} \rrbracket^C = \lambda P \lambda x. \exists d [\mathbf{standard}(d, P, C) \wedge P(d)(x)]$
 b. $\llbracket \text{believes } S \rrbracket = \lambda d \lambda x. Cr_x(\llbracket S \rrbracket) \geq d$
 c. $\llbracket [\text{pos}_{Adv} \text{ strongly}] [\text{believes } S] \rrbracket = \lambda x. \exists d [d > d_s \wedge Cr_x(\llbracket S \rrbracket) \geq d]$
 $= \lambda x. Cr_x(\llbracket S \rrbracket) > d_s$

A standard meaning for adjectival *more* is given in (35a); here d' is filled by the degree denoted by the comparative clause (Kennedy, 1999). The adverbial counterpart in (35b) makes room for an additional argument supplied by the gradable adverb.

- (35) a. $\llbracket \text{more} \rrbracket = \lambda d' \lambda P \lambda x. \exists d [P(d)(x) \wedge d > d']$
 b. $\llbracket \text{more}_{Adv} \rrbracket = \lambda A \lambda d' \lambda P \lambda x. \exists d [A(P)(d)(x) \wedge d > d']$

Assuming the elliptical structure in (36), we can derive comparative uses involving *believe* as follows.

- (36) John $\llbracket \text{believes } S \rrbracket \llbracket [\text{more}_{Adv} \text{ strongly}] \llbracket \text{than Mary believes } S \rrbracket \rrbracket$
 a. $\llbracket \text{more}_{Adv} \text{ strongly} \rrbracket = \lambda d' \lambda P \lambda x. \exists d [P(d)(x) \wedge d > d']$
 b. $\llbracket \text{than Mary believes } S \rrbracket = d_c$
 c. $\llbracket [\text{more}_{Adv} \text{ strongly}] \llbracket \text{than Mary believes } S \rrbracket \rrbracket = \lambda P \lambda x. \exists d [P(d)(x) \wedge d > d_c]$
 d. $\llbracket \llbracket \text{believes } S \rrbracket \llbracket [\text{more}_{Adv} \text{ strongly}] \llbracket \text{than Mary believes } S \rrbracket \rrbracket \rrbracket$
 $= \lambda x. \exists d [Cr_x(\llbracket S \rrbracket) \geq d \wedge d > d_c]$
 $= \lambda x. Cr_x(\llbracket S \rrbracket) > d_c$

This should suffice to demonstrate the general plausibility of treating *believe* as a gradable verb. Needless to say, what has been presented in this subsection is just an outline of a rich and very complex phenomenon. But it is promising that it already provides the basic functionality and is well incorporated into the degree-based framework of gradability.

3.2. Conjunction closure

One desirable consequence of the current account is that *believe* is closed under conjunction. To see that, assume that *x believes p* and *x believes q* are both true. Given the proposal developed in the previous subsection, it follows that $Cr_x(p) = 1$ and $Cr_x(q) = 1$. These statements say that, according to *x*, the entire probability weight falls within *p* and the entire probability weight falls within *q*. That is, according to *x*, all possible worlds outside *p* and outside *q* are assigned a probability of zero. This entails that $Cr_x(p \cap q) = 1$, i.e., *x believes (p and q)* is true. Given the empirical data discussed in subsection 2.3, this is a welcome result.

3.3. Explaining hedging

If *believe* entails full confidence in the prejacent, as the bulk of the empirical evidence presented in section 2 suggests, how can we reconcile this with the possibility of hedging? I argue that the intuition of weakness falls out as a particular kind of scalar inference that arises through competition with a presuppositionally stronger *know*-alternative. In order to understand why this inference to weakness does not clash with the full strength of *believe*, we need to draw a distinction between two types of certainty, i.e., subjective and objective. **Subjective** certainty is what is expressed in belief reports; it is privately held and need not be based on empirical evidence. By contrast, **objective** certainty is publicly expressed and needs to be backed up by evidence. This is the kind of certainty that is at stake in conversation, as it entails commitments and carries with it the burden of proof. Subjective certainty is measured by the probability function *Cr*, already introduced above. Objective certainty, I assume, is measured by the function *Pr*. One can think of *Pr* as a more conservative version of *Cr*, although this will only hold for sincere speakers. That is, if a speaker is publicly committed to a proposition to a certain degree, her subjective confidence in that proposition will normally meet that degree: $Pr_x(p) \leq Cr_x(p)$,

for all sincere agents x and propositions p .⁶

I now discuss the nature of the scalar inference that is responsible for the intuition of weakness. Hawkins (1991), following along the classical analysis of Russell (1905), proposed that *the* and $a(n)$ share an entailment of existence but *the* additionally introduces an implication of uniqueness. Hawkins assumed that this uniqueness implication is a regular entailment, so that $\langle a(n), the \rangle$ constitutes an entailment-based lexical scale or an **e-scale**. The inference to non-uniqueness associated with $a(n)$ is then derived as a scalar implicature by standard neo-Gricean reasoning. However, Heim (1991) argued that the uniqueness implication of the definite article is a presupposition rather than an entailment. The non-uniqueness inference associated with the indefinite article would then arise through competition with a lexical item that is presuppositionally stronger, and hence this inference cannot be a regular implicature. Heim proposed that it is derived by the principle of **Maximize Presupposition**, which states that among two sentences with (contextually) equivalent truth conditions the one with the stronger presupposition is to be preferred, provided that these presuppositions are met. This principle explains why, for example, we cannot felicitously utter *A sun is shining*. The speaker should rather utter *The sun is shining*, given that in our solar system there is a single sun and thus the presupposition of *the* is satisfied.

Later work has added to $\langle a(n), the \rangle$ more instances of presupposition-based scales, or **p-scales**, including $\langle all, both \rangle$, $\langle PL, SING \rangle$, $\langle PRES, PAST \rangle$, $\langle \emptyset, too \rangle$, $\langle believe, know \rangle$ (Percus, 2006; Chemla, 2008; Sauerland, 2008; Singh, 2011; Schlenker, 2012). In this paper, I am interested in the p-scale $\langle believe, know \rangle$, where it is assumed that its elements share (contextually) equivalent truth conditions, but *know* adds a presupposition to the effect that its complement is true. Maximize Presupposition then helps us understand why the use of *believe* implies that the presupposition of the *know*-alternative is not certain to hold. Following Percus (2006), I call this type of scalar inference an **antipresupposition**.⁷

- (37) a. Actual utterance: John believes it's raining.
 b. Alternative utterance: John knows it's raining.
 c. Antipresupposition: It's not certain that it's raining.

The following lexical entry views *know* as a factive counterpart of *believe*. Below, I add possible worlds to the metalanguage and adopt Heim and Kratzer's (1998) convention of placing presuppositions between the lambda operators and the scope.

$$(38) \quad \llbracket \text{know} \rrbracket^w = \lambda p \lambda d \lambda x : p(w) \cdot Cr_{x,w}(p) \geq d$$

This entry entails that *know* is gradable, and this is debatable (for discussion, see Partee, 2004; Stanley, 2005). In the face of felicitous examples like *He knows very/quite/full well that I don't like alcohol*, I will tentatively assume that *know* does make available a degree argument, as in (38). Nothing important depends on this choice, though. The above entry can easily be modified to a non-gradable one as follows: $\llbracket \text{know} \rrbracket^w = \lambda p \lambda x : p(w) \cdot Cr_{x,w}(p) = 1$. The above entry also

⁶The way I use the terms “subjective” and “objective” does not quite line up with philosophical parlance on probability. There are several interpretations of probability, including frequentist (probability as chance or proportion) and subjective or Bayesian (probability as a measure of an agent's certainty); see Hájek (2011) for an overview. *Cr* and *Pr* are both “subjective” in the broad sense of being tied to an agent.

⁷To be precise, Percus actually calls antipresupposition what is not taken to hold; in (37), that would be the proposition that it is raining. This use of the term is less common and I will not adopt it here.

entails that *know* is truth-conditionally equivalent to *believe* (for qualifying this to “contextual” equivalence, see Chemla, 2008; Schlenker, 2012).

We can now derive the antipresupposition of *believe* through competition with *know*. I propose to do this by means of a presupposition-based exhaustivity operator that captures the essence of Maximize Presupposition but has the advantage that it may occur in subordinate clauses and generate antipresuppositions locally (cf. *Mike said he believes Kamala will win but he is not sure*).

Let $\langle \alpha_1, \dots, \alpha_n \rangle$ be a p-scale and $S(\alpha_i)$ be a Logical Form that contains α_i , an element of this scale. We can define the set of presuppositional alternatives, or **p-alternatives** of $S(\alpha_i)$, as the set of all structurally similar Logical Forms in which α_i may be substituted by one of its scale-mates.

$$(39) \quad p\text{-Alt}(S(\alpha)) = \{S(\beta) \mid \alpha \text{ and } \beta \text{ belong to the same p-scale}\}$$

Only p-alternatives with a stronger presupposition than the uttered sentence are excludable, i.e., can be denied in order to obtain an enriched meaning. The set of excludable p-alternatives is defined below, where ∂ is a presupposition operator (adapted from Beaver, 2001) that isolates the presupposition of a sentence meaning.

$$(40) \quad p\text{-Excl}(S) = \{S' \in p\text{-Alt}(S) \mid \partial[S'] \subset \partial[S]\} \quad (\text{first version})$$

Work on embedded scalar implicatures has employed a silent exhaustivity operator *Exh* that attaches to a clause and enriches its meaning with the condition that all excludable scalar alternatives are false (Chierchia, 2008, 2013; Fox, 2007; Chierchia et al., 2012). I introduce a presuppositional counterpart *p-Exh*, which adds the condition that the speaker (marked as *s*) is not objectively certain that the presupposition of any stronger p-alternative holds.⁸ This essentially means that the speaker lacks appropriate evidence that any such presupposition is true.

$$(41) \quad \llbracket p\text{-Exh } S \rrbracket^w = \llbracket S \rrbracket^w \wedge \forall S' \in p\text{-Excl}(S) : Pr_{s,w}(\partial[S']) < 1 \quad (\text{first version})$$

Here is how *p-Exh* derives the antipresupposition associated with *believe*-sentences. The final line below states that John is subjectively certain it is raining but the speaker lacks appropriate evidence for this being the case.

$$(42) \quad p\text{-Exh} [\text{John } pos \text{ believes it's raining}]$$

- a. $\llbracket \text{John } pos \text{ believes it's raining} \rrbracket^w = Cr_{[\text{John}],w}(\llbracket \text{rain} \rrbracket) = 1$
- b. $p\text{-Alt}(\text{John } pos \text{ believes it's raining}) = \left\{ \begin{array}{l} \text{John } pos \text{ believes it's raining,} \\ \text{John } pos \text{ knows it's raining} \end{array} \right\}$
- c. $p\text{-Excl}(\text{John } pos \text{ believes it's raining}) = \{\text{John } pos \text{ knows it's raining}\}$
- d. $\llbracket p\text{-Exh} [\text{John } pos \text{ believes it's raining}] \rrbracket^w$
 $= Cr_{[\text{John}],w}(\llbracket \text{rain} \rrbracket) = 1 \wedge Pr_{s,w}(\llbracket \text{rain} \rrbracket) < 1$

Chemla (2008) notices that the antipresupposition of *believe* can be invoked to explain the intuition of weakness with first-person belief attributions. Indeed, the structure in (43) derives

⁸The exact way in which the exclusivity component is framed in the literature on antipresupposition varies depending on the modality involved (belief, knowledge, authority), the responsible agent (the speaker or all discourse participants), and the scope of the negation (wide or narrow with respect to the modal operator). The specific choices do not matter to our purposes as long as what is denied is the speaker's objective certainty.

the hedging use by following the same steps as in (42). We get an interpretation according to which the speaker is subjectively certain but is not committed to it being the case that it is raining, presumably because she lacks sufficient evidence to back up her claim.

- (43) a. *p-Exh* [I *pos* believe it's raining]
 b. $Cr_{s,w}(\llbracket \text{rain} \rrbracket) = 1 \wedge Pr_{s,w}(\llbracket \text{rain} \rrbracket) < 1$

The truth condition produced above makes it clear why the hedging use of *x* believes that *p* is naturally spelled out by a follow-up clause along the lines of *x* is not certain that *p*. If we agree that *certain* encodes objective probabilities (e.g., $\llbracket \text{certain} \rrbracket^w = \lambda p \lambda d \lambda x. Pr_{x,w}(p) \geq d$), its negation will have the same effect as the exclusivity inference triggered by *p-Exh*.

Our final task is to derive the sensitivity of hedging to the question under discussion. Recall from (24)-(25) that the intuition of weakness arises only when the belief attribution itself is not relevant to the question under discussion. This pattern can be explained if we assume that relevant *p*-alternatives are filtered out by the computational system. Why should relevant (rather than non-relevant) *p*-alternatives be excluded by the system? The rationale behind this assumption is that *p*-alternatives (as the name suggests) are based on presuppositions, and these are typically not relevant.⁹ I follow the bulk of the literature in assuming that a proposition is **relevant** to a question if it provides a partial answer to that question, i.e., if it is incompatible with at least one possible answer (Groenendijk and Stokhof, 1984; Roberts, 1996; van Rooy, 2003; Simons et al., 2010).

- (44) $\mathbf{rel}(p, Q)$ iff $\exists q \in Q : p \cap q = \emptyset$

We can now restrict the set of *p*-alternatives that underlies the exhaustivity operator to non-relevant propositions.

- (45) $p\text{-Excl}_Q(S) = \{S' \in p\text{-Alt}(S) \mid \partial \llbracket S' \rrbracket \subset \partial \llbracket S \rrbracket \wedge \neg \mathbf{rel}(\llbracket S' \rrbracket, Q)\}$ (final version)

Given this refinement, we can understand why weakness arises only if the excludable alternatives are not relevant. This is illustrated schematically by the following two examples. In (46), the *know*-alternative is not relevant to the question under discussion: given that relevance only cares about truth-conditional content, the fact that an agent assigns maximal credence to a proposition does not decide on its truth. Thus, this alternative survives and we correctly predict that a hedging use is available. In (47), by contrast, the *know*-alternative is relevant (it is incompatible with the second question alternative) and does not survive. As a result, exhaustification has no semantic effect and a hedging use is not available.

- (46) a. {Is global warming real?} $Q = \{r, \neg r\}$
 b. *p-Exh* [I believe global warming is real] $p\text{-Alt} = \{B_s r, K_s r\}$, $p\text{-Excl}_Q = \{K_s r\}$
 c. Enriched meaning: $B_s r \wedge \neg K_s r$

⁹This fact lends further support to the claim that the inference to weakness is an antipresupposition rather than a scalar implicature. Scalar implicatures are based on entailment and exhibit the reverse pattern, i.e., they typically arise only if relevant (Romoli, 2013). Compare (i), where the implicature is relevant and difficult to cancel, to (ii), where the implicature is not relevant and very easy to cancel.

- (i) Q: Did you read the articles the professor recommended?
 A: I read some of them. ?In fact, I read all of them.
 (ii) Q: Why did you remove the first slide from your class presentation?
 A: Some of the students found it offensive. In fact, all of them did.

- (47) a. {Tell us about your environmental beliefs.} $Q = \{B_{sr}, B_s \neg r, \dots\}$
 b. $p\text{-Exh}$ [I believe global warming is real] $p\text{-Alt} = \{B_{sr}, K_{sr}\}, p\text{-Excl}_Q = \emptyset$
 c. Enriched meaning (=basic meaning): B_{sr}

The final version of the semantics for the p-exhaustivity operator is catalogued below.

- (48) $\llbracket p\text{-Exh } S \rrbracket^{c,w} = \llbracket S \rrbracket^{c,w} \wedge \forall S' \in p\text{-Excl}_Q(S) : Pr_{s,w}(\partial \llbracket S' \rrbracket^c) < 1$ (final version)

4. Conclusion

The idea that *believe* expresses universal quantification over possible worlds hails from a long and venerable tradition in formal semantics. This paper challenged this mantra as both too rigid and non-explanatory, pointing out that it does not capture the gradability of *believe* and fails to predict that (an unmodified use of) this verb entails full subjective certainty. I have argued for a semantics that views *believe* as a maximum-degree predicate, allows for its strength to be manipulated by degree modifiers, and correctly predicts that it is closed under conjunction. Importantly, I assumed that the probability measure encoded by *believe* is subjective, which in the right context can give rise to the intuition of weakness. I have shown that this intuition can be construed (and appropriately constrained) as a scalar inference due to the presence of a covert exhaustivity operator that compares alternatives of different presuppositional strength.

References

- Bartsch, R. (1973). ‘Negative transportation’ gibt es nicht. *Linguistische Berichte* 27, 1–7.
- Beaver, D. I. (2001). *Presupposition and Assertion in Dynamic Semantics*. Stanford: Center for the Study of Language and Information.
- Büring, D. (2003). On d-trees, beans, and b-accents. *Linguistics and Philosophy* 26(2), 511–545.
- Chemla, E. (2008). An epistemic step for anti-presuppositions. *Journal of Semantics* 25, 141–173.
- Chierchia, G. (2008). An epistemic step for anti-presuppositions. *Linguistic Inquiry* 37(4), 535–590.
- Chierchia, G. (2013). *Logic in Grammar: Polarity, Free Choice, and Intervention*. Oxford: Oxford University Press.
- Chierchia, G., D. Fox, and B. Spector (2012). Scalar implicature as a grammatical phenomenon. In C. Maienborn, K. von Stechow, and P. Portner (Eds.), *Semantics: An International Handbook of Natural Language Meaning*, pp. 2297–2331. Berlin/Boston: De Gruyter Mouton.
- Clark, R. (2013). Belief is credence one (in context). *Philosophers’ Imprint* 13(11), 1–18.
- Cresswell, M. J. (1976). The semantics of degree. In B. Partee (Ed.), *Montague Grammar*, pp. 261–292. New York: Academic Press.
- Farkas, D. and K. B. Bruce (2010). On reacting to assertions and polar questions. *Journal of Semantics* 27, 81–118.
- Foley, R. (1992). The epistemology of belief and the epistemology of degrees of belief. *American Philosophical Quarterly* 29(2), 111–124.
- Fox, D. (2007). Free choice and the theory of scalar implicatures. In U. Sauerland and P. Stat-

- eva (Eds.), *Presupposition and Implicature in Compositional Semantics*, pp. 71–120. Basingstoke: Palgrave.
- Gajewski, J. R. (2007). Neg-raising and polarity. *Linguistics and Philosophy* 30, 289–328.
- Ginzburg, J. (1996). Interrogatives: Questions, facts, and dialogue. In S. Lappin (Ed.), *The Handbook of Contemporary Semantic Theory*, pp. 385–422. Oxford: Blackwell.
- Ginzburg, J. (2012). *The Interactive Stance*. Oxford: Oxford University Press.
- 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.
- Hájek, A. (2011). Interpretations of probability. In *Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/entries/probability-interpret/>.
- Hawkins, J. (1991). On (in)definite articles: implicatures and (un)grammaticality prediction. *Journal of Linguistics* 27, 405–442.
- Hawthorne, J., D. Rothschild, and L. Spectre (2016). Belief is weak. *Philosophical Studies* 173, 1393–1404.
- Heim, I. (1991). Artikel und Definitheit. In A. v. Stechow and D. Wunderlich (Eds.), *Semantik: Ein internationales Handbuch der zeitgenössischen Forschung*, pp. 487–535. Berlin: de Gruyter.
- Heim, I. and A. Kratzer (1998). *Semantics in Generative Grammar*. Oxford: Blackwell.
- Hintikka, J. (1962). *Knowledge and Belief*. Ithaca: Cornell University Press.
- Hintikka, J. (1969). Semantics for propositional attitudes. In J. W. D. et al. (Ed.), *Philosophical Logic*, pp. 21–45. Dordrecht: Reidel.
- Homer, V. (2015). Neg-raising and positive polarity: The view from modals. *Semantics & Pragmatics* 8(4), 1–88.
- Horn, L. R. (1989). *A Natural History of Negation*. Chicago: University of Chicago Press.
- Kennedy, C. (1999). *Projecting the Adjective: The Syntax and Semantics of Gradability and Comparison*. Garland Press: New York.
- Kennedy, C. (2007). Vagueness and grammar: The semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy* 30(1), 1–45.
- Kennedy, C. and L. McNally (2005). Scale structure, degree modification, and the semantics of gradable predicates. *Language* 81(2), 345–381.
- Klecha, P. (2014). *Bridging the Divide: Scalarity and Modality*. University of Chicago: Unpublished dissertation.
- Kratzer, A. (1991). Modality. In A. von Stechow and D. Wunderlich (Eds.), *Semantik: Ein internationales Handbuch der zeitgenössischen Forschung*, pp. 639–650. de Gruyter.
- Kyburg, H. E. (1961). *Probability and the Logic of Rational Belief*. Middletown, CT: Wesleyan University Press.
- Laserson, P. (2005). Context dependence, disagreement, and predicates of personal taste. *Linguistics and Philosophy* 28, 643–686.
- Lassiter, D. (2017). *Graded Modality: Qualitative and Quantitative Perspectives*. Oxford: Oxford University Press.
- Lauer, S. (2017). ‘I believe’ in a ranking-theoretic analysis of ‘believe’. In *Proceedings of the Twenty-first Amsterdam Colloquium*.
- Leitgeb, H. (2014). The stability theory of belief. *Philosophical Review* 123(2), 173–204.
- Levi, I. (1973). *Gambling With Truth: An Essay On Induction and the Aims of Science*. Cam-

- bridge, MA: MIT Press.
- Levi, I. (1991). *The Fixation of Belief and Its Undoing*. Cambridge, MA: Cambridge University Press.
- Partee, B. (2004). Comments on Jason Stanley's 'on the linguistic basis for contextualism'. *Philosophical Studies* 119(1–2), 147–159.
- Percus, O. (2006). Antipresuppositions. In A. Ueyama (Ed.), *Theoretical and Empirical Studies of Reference and Anaphora: Toward the establishment of generative grammar as an empirical science*, pp. 52–73. Report of the Grant-in-Aid for Scientific Research (B), Project No. 15320052, Japan Society for the Promotion of Science.
- Portner, P. (2009). *Modality*. Oxford: Oxford University Press.
- Roberts, C. (1996). Information structure: Towards an integrated formal theory of pragmatics. In J. Yoon and A. Kathol (Eds.), *OSU Working Papers in Linguistics 49: Papers in Semantics*, pp. 91–136.
- Romoli, J. (2013). A scalar implicature-based approach to neg-raising. *Linguistics and Philosophy* 36, 291–353.
- Russell, B. (1905). On denoting. *Mind* 14(56), 479–493.
- Santorio, P. and J. Romoli (2017). Probability and implicatures: A unified account of the scalar effects of disjunction under modals. *Semantics & Pragmatics* 10(13).
- Sauerland, U. (2008). Implicated presuppositions. In A. Steube (Ed.), *The Discourse Potential of Underspecified Structures*, pp. 581–600. Berlin: Mouton de Gruyter.
- Schlenker, P. (2012). *Maximize Presupposition* and Gricean reasoning. *Natural Language Semantics* 20, 391–429.
- 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.
- Singh, R. (2011). *Maximize Presupposition!* and local contexts. *Natural Language Semantics* 19, 149–168.
- Stanley, J. (2005). *Knowledge and Practical Interests*. New York/Oxford: Oxford University Press.
- Unger, P. (1971). A defense of skepticism. *The Philosophical Review* 80(2), 198–219.
- van Rooy, R. (2003). Questioning to resolve decision problems. *Linguistics and Philosophy* 26(6), 727–763.
- von Stechow, K. and A. Gillies (2010). *Must ... stay ... strong!* *Natural Language Semantics* 18(4), 351–383.
- von Stechow, A. (1984). Comparing semantic theories of comparison. *Journal of Semantics* 3, 1–77.
- Williamson, T. (2002). *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Yalcin, S. (2010). Probability operators. *Philosophy Compass* 5(11), 916–937.