Hedging verbs and nouns using an alternative semantics

Curt Anderson

In this paper I provide a formal analysis of the English hedge *sorta*, concentrating on its use with verb phrases. I bring to light new data showing how *sorta* can hedge the direct object of creation verbs and some intensional verbs without combining with the direct object directly, and that the ability to hedge a direct object from a distance is conditioned by verb type and the type of determiner with the direct object. I build an analysis using Morzycki's (2011) alternative semantics implementation of the pragmatic halos of Lasersohn (1999).

1. Introduction

Speech is rarely perfectly precise. Well-known from the work of Grice (1957) is that speakers are experts at implicated unsaid meanings — how something is said matters to the way it is interpreted. But the competence system allows for other avenues to express things that are not quite what was said; quantifiers can have restricted domains, and hedges such as *loosely speaking* and *sorta* mark words and expressions as being interpreted in a way that is outside the norm. This paper is an analysis of one such hedge, *sorta*.

The question pursued in this paper is about the representation behind hedging, and how *sorta* can grade over different meanings. I show that *sorta* is sensitive to noun phrase and verbal semantics, and that accounting for its behavior requires new assumptions about the denotations of lexical items. In doing this, I adopt a Hamblin semantics for imprecision from Morzycki (2011), who suggests modeling the denotations of lexical items as sets of alternatives, this being a reimplementation of Lasersohn (1999)'s pragmatic halos.

This paper starts by laying out the crucial data in section 2. From there, I move on to discussing the theoretical background of this paper, pragmatic halos and a Hamblin semantics for imprecision, in section 3. A first pass at an analysis is provided in section 4, with further discussion of the analysis in section 5. I discuss implications for the interface between semantics and the conceptual system in section 6, as well as a way forward in tackling the puzzles that confound my analysis.

Proceedings of ConSOLE XXI, 2013, 1-20
http://www.sole.leidenuniv.nl

2. The basic data

2.1. Approximation and sorta

The English modifier *sorta* is part of a family of modifiers often called *hedges*.[1] Like other hedges, *sorta* pragmatically serves to signal a mismatch between what a speaker is saying and what a speaker actually means. This may be done for the metalinguistic reason of not knowing the correct word or phrase to use at the time of utterance, or as a rhetorical strategy to soften the impact of what is said. In this respect, it behaves similar to *like* (Siegel 2002).

The mismatch *sorta* provides between what is said and meant is demonstrated in (1) below, where the natural interpretation is that the speaker does not mean to actually use the word *kick* in the sentence, but some other word. Denying the use of the word *kick* in a followup is perfectly acceptable (2), as the semantic content of *kick* is neutralized by *sorta* (Bolinger 1972).

(1) [I was] running on concrete and accidentally sorta kicked the ground. (Google)

(2) I was running on concrete and accidentally sorta kicked the ground — that is to say, I didn't really kick the ground, but it was like kicking the ground.

The approximative sense with *sorta* comes out clearly when considering the paraphrases available when *sorta* is used. The most natural paraphrases are those that express that *sorta V* (where V is the verb) is similar to but not V in some respect.

(3) He sorta swam over to the boat.
"He did something like swimming over to the boat."

(4) The soccer player sorta kicked the ball.
"The player did something close to but not quite kicking the ball."

Sorta is able to modify verb phrases headed by most types of verbs. For many verbs, the behavior of *sorta* mirrors the behavior in (1) above, where *sorta* hedges the verb. For some verbs, however, a second reading arises where *sorta* can hedge not just the verb but also its direct object. Verbs that *sorta* can do this for include some intensional transitive verbs (such as *look for*), as well as creation verbs (*build*), depiction verbs (*draw*, *paint*), and performance verbs (*sing*). To illustrate, in (5) below, the sentence is ambiguous between two readings: one reading where the verb is hedged, but also one reading where the direct object of the verb is hedged. (6) demonstrates the same phenomenon with a depiction verb, and (7) with *look for*.

(5) The carpenter sorta built a barn.
   a. The carpenter did something that was like building a barn (e.g., putting together a prefabricated structure).

[1] In this paper I am concerned with the adverbial *sorta* that can modify VPs and APs. The nominal *sort of*, i.e. *a sort of cat*, does not figure into this story. The adverbial *sorta* that can modify VPs and APs has a phonologically unreduced variant *sort of*, but to emphasize my treating adverbial *sort of* and *sorta* as a single lexical item separate from the nominal *sort of*, I write both as *sorta*, even when the unreduced variant is the preferred variant for speakers. *Kind of* (reduced: *kinda*) appears to be related to *sorta*, and for the purposes of this paper I assume that it has the same semantics and pragmatics. Some speakers report that judgements with *sorta* are odd while using *kinda* is better, but I assume that this is stylistic variation.

b. The carpenter built something like a barn (e.g., a shed).

(6) The boy sorta drew a house.

a. The boy did something like drawing a house (e.g., connected the dots in a picture).

b. The boy drew something that was like a house.

(7) I'm sorta looking for a horse.

a. I'm only half-heartedly looking for a horse.

b. I'm looking for something like a horse.

Looking at the examples above, we might think that this is simply an effect of the indefiniteness of the direct object. To some extent, it is; replacing the indefinite noun phrases with definite noun phrases in the sentences above makes the hedging effect on the direct object disappear. Both (8) and (9) below are unambiguous, showing that the determiner does have an effect.

(8) The amateur carpenter sorta built the house.

a. The amateur carpenter did something that was akin to building that resulted in the house (e.g., he had help from more skilled carpenters).

b. *The amateur carpenter built something that was like the house (a shack, a hovel, . . .).

(9) The men sorta sang the song.

a. The men did something that was like singing (mumbling, bellowing, . . .).

b. *The men sang something that was like the song (a poem, a verse, . . .).

When we look at more mundane verbs such as *kick* with indefinite direct objects, however, a different picture begins to emerge. (10) does not have the same sorts of readings that verbs such as *build* above do, and neither does *eat* in (11). Namely, the reading where the direct object is hedged is unavailable, even though the verb may still be hedged. Both of these verbs are outside of the verb classes delineated above, the creation verb and intensional verb classes.

(10) The soccer player sorta kicked a ball.

a. The soccer player did something that was like kicking to a ball.

b. *The soccer player kicked something that was like a ball.

(11) The woman sorta ate a cracker.

a. The woman did something that was like eating to a cracker.

b. *The woman ate something that was like a cracker.

What should be clear from this data is that there is a confluence of both indefinite noun phrase direct objects and verb class in getting this additional reading with *sorta*. First, only indefinite objects may be hedged — definite objects resist the hedging effects of *sorta*. Second, indefinite objects may only be hedged if they are selected for by one of a particular class of verbs, namely verbs of creation or intensional verbs.

A matter worth reflecting on for this puzzle is how *sorta* can affect a direct object at all. Given standard syntactic assumptions, *sorta* and the direct object never form a constituent, and in fact, *sorta* directly modifying a noun phrase is quite degraded (12, 13).[2]

(12) *I saw sorta a bird.

(13) *She ate sorta a cracker.

Under common assumptions about compositionality, it should be a bit of a mystery about how *sorta* can affect the interpretation of an NP when it does not form a syntactic constituent with it.

2.2. Summary

When hedging verbs, some but not all verbs allow for their direct objects to be hedged by *sorta*. The relevant factor seems to be verb class. Complicating this generalization is that, for those verbs that allow their direct object to be hedged, it is only indefinite objects that are allowed to be hedged. Definite direct objects are never allowed to be hedged. This naturally leads to several questions about the nature of *sorta*. First, how can we represent the effect that *sorta* has on words, the so-called hedging effect? Second, what representation of *sorta* can derive the correct behavior of *sorta* with certain types of objects — why can indefinite objects but not definite objects be hedged? Finally, why is it that the indefinite objects of certain verbs are special? Why are only indefinite objects available for hedging with *sorta*?

3. Approximation

3.1. Pragmatic halos

Examining natural language expressions, Lasersohn (1999) notes that the pragmatics of many expressions allows them to be used even in situations where they would normally be considered false, strictly speaking. For example, consider the context and discourse in (14).

[2]Assuming a syntax where a D(eterminer)P(hrase) takes an NP complement (Abney 1987), *sorta* can modify the NP in limited ways for some speakers. The adverbial *sorta* in this case contrasts with the noun *sort*, as demonstrated below. For instance, *a sort of fairytale* is a type of fairytale, but *a sorta fairytale* is something like a fairytale. Likewise, a Porsche is a sort of car, but not a sorta car, since it is definitely a car.

(1) a. a sort of fairytale
“a type of fairytale”

b. a sorta fairytale
“almost but not a fairytale”

(2) a. A Porsche is a sort of car.

b. #A Porsche is a sorta car.

My analysis seems consistent with these facts about noun phrases, but I do not pursue any further analysis of the DP-internal adverbial *sorta* in this paper.

(14) *Jen is discussing with Jim and Jill when various people arrived at their party the night before. Jack arrived at exactly 6:58pm.*
Jen: What time did Jack arrive?
Jim: He arrived at 7pm.

(15) #**Jill**: No, he arrived at 6:58pm.

In this discourse, we do not consider Jim to have said anything particularly pathological in terms of discourse structure. What Jim said was perfectly licit, given the standards of the conversation. In light of these standards, it is in fact odd for Jill to follow Jim's comment with the comment in (15). This is even more puzzling considering the facts of the matter: Jack did in fact arrive at 6:58pm. What Jim has said is absolutely false given this fact, and what Jill has said is absolutely true, but Jill's comment is illicit with respect to the conversation at hand while Jim's is perfectly acceptable.

Lasersohn argues that discourses allow for a certain amount of pragmatic slack or imprecision in interpreting expressions. Although certain things people say may in fact be absolutely false, discourses allows for an amount of leeway in what expressions count as good enough to be used. In casual speech, speakers typically allow each other quite a bit of pragmatic slack, as shown in the fictional discourse above. Jim's statement that Jack arrived at 7pm is good enough given the aims of the conversation, even though it is not truthful in the strictest sense. Jill's follow-up comment, although true, is regarded as odd because it is too precise given the imprecision allowed in this context. In fact, such utterances can be accommodated so long as we acknowledge we are entering a context where such pedantry is tolerated.

(16) **Jim**: Well, I guess that's true, Jill, but you're just being a pedant.

Lasersohn suggests that an appropriate way to model the effect of imprecision in discourse is to consider natural language expressions as projecting two types of meaning. One type of meaning is the ordinary truth conditional meaning of an expression, the other being a set of pragmatically ignorable differences given the context that he calls a pragmatic halo. Speakers are allowed imprecision in the meaning of their words and phrases so long as they fall within the pragmatic halo of an expression. For the discourse above, Jen and Jim will still interpret *7pm* as **7pm**, but implicitly acknowledge that *6:58pm* falls within the pragmatic halo of **7pm** and is hence an ignorable difference given the imprecision allowed to each other.

Not only do pragmatic halos shrink or expand implicitly with the context, but certain expressions that Lasersohn terms slack regulators can also affect the size of pragmatic halos. *Exactly* is argued to be a slack regulator, having the ability to shrink the size of a halo and hence allowing for less imprecision in how an expression may be evaluated. The effect of this can be seen in the discourse below, where Jen asks the exact time that Jack arrived.

(17) *Jen is discussing with Jim and Jill when various people arrived at their party the night before. Jack arrived at exactly 6:58pm.*
Jen: At exactly what time did Jack arrive?
*__Jim__: He arrived at 7pm.
Jill: No, he arrived at 6:58pm.

In this discourse, Jim's statement is false as before, but it is now also pragmatically unacceptable, due to Jen's insistence on knowing the exact time that Jack arrived. The effect of *exactly* here was to shrink the pragmatic halo enough that *6:58pm* was excluded from the halo. Jill's follow-up that Jim is wrong is still true in this discourse, but now also pragmatically acceptable, given the fact that Jen has required more precision in the discourse by virtue of using *exactly*. And, Jim's response is now unacceptable, due to the increased amount of precision.

I suggest that *sorta* has some of the same flavor as a slack regulator. Namely, what *sorta* does is allow for normally ignorable expressions to be considered in place of the expression that *sorta* modifies. If *sorta* α is some expression, *sorta* allows for variants to α that would not otherwise be available given the discourse. This is essentially a slack regulating function, but with an important difference — although slack regulators like *exactly* shrink the size of the pragmatic halo, *sorta* does the opposite in increasing the size of the pragmatic halo.

3.2. An alternative semantics for halos

Morzycki (2011) analyzes metalinguistic comparatives (McCawley 1998; Giannakidou & Stavrou 2009; Giannakidou & Yoon 2011) as indirectly comparing the size of pragmatic halos. Morzycki assumes that the interpretation function $[\![.]\!]$ is parameterized to a degree of precision d (which I call the degree of precision or imprecision parameter), similar to how the interpretation function can be parameterized to a world in an intensional system. The degree of precision is directly related to the pragmatic halo for a given expression. A simplified denotation for (18) is given in (19).

(18) George is more dumb than crazy.

(19) $[\![\textit{George is more dumb than crazy}]\!]^{d'}$
$= \textbf{max}(\lambda d.[\![\textit{George is dumb}]\!]^{d}) \succ \textbf{max}(\lambda d.[\![\textit{George is crazy}]\!]^{d})$

What the metalinguistic *more* does in this case is compare the degree to which George can be called dumb with the degree to which he can be called crazy.

Pragmatic halos in Morzycki's analysis are modeled as Hamblin alternatives (Hamblin 1973; Rooth 1985, 1992; Kratzer & Shimoyama 2002). These halos are functions that resemble the core predicate denoted by the lexical item to some degree. For instance, the halo for the adjective *dumb* might include not just the function **dumb**, which we might take *dumb* to standardly denote, but also the functions **dopey**, **foolish**, and so on, given an appropriate amount of pragmatic slack.

Halos are generated as sets of resembling functions. To accomplish this, Morzycki introduces a new relation, $\approx$ "resembles", which is true just in case two objects resemble each other to at least degree d (see 20). d is a degree in the real interval $[0, 1]$, and when $d = 1$, $\approx$ is formally equivalent to $=$.

(20) $\alpha \approx_{d,C} \beta$ iff, given the ordering imposed by the context C, α resembles β to (at least) the degree d and α and β are of the same type (Morzycki 2011).

Denotations are conceived of as their pragmatic halos, with the degree of precision parameter d on the interpretation function controlling the size of the set of alternatives. As the degree in-

creases towards the maximal degree, the halo contracts, while the halo expands as d approaches the minimum degree. Putting this together, *dumb* might be represented as in (21).

(21) $[\![\textit{dumb}]\!]^d = \{f : f \approx_{d,C} \textbf{dumb}\}$

Morzycki provides a typeshift PREC which can be used to get access to the degree of precision parameter (22). The content of PREC is to simply abstract over the degree of precision. This has the effect of turning any expression type $\langle\tau\rangle$ into type $\langle d, \tau\rangle$.

(22) $[\![\text{PREC } \alpha]\!]^d = \lambda d'.[\![\alpha]\!]^{d'}$

(23) $[\![\text{PREC } \textit{dumb}]\!]^d = \lambda d'.[\![\textit{dumb}]\!]^{d'}$

In a Hamblinized system such as this, function application cannot proceed per the usual (e.g., Heim & Kratzer (1998)'s FUNCTION APPLICATION (FA)), due to denotations being sets rather than functions.[3] But, even though these are sets, we would like to think of denotations as having the same type as the objects in their sets of alternatives. For example, although $[\![\textit{dumb}]\!]^d$ is the set of objects resembling **dumb** to degree d, we would still like to think of this as being type $\langle e, st\rangle$. Therefore, we need a new notion of what it means to apply one expression to another. The intuition is to apply all the objects from one set of alternatives to all the objects from another set of alternatives pointwise, creating another set of alternatives. This is formalized as HAMBLIN FUNCTION APPLICATION in (24) below.

(24) HAMBLIN FUNCTIONAL APPLICATION (HFA)
If α is a branching node with daughters β and γ, and $[\![\beta]\!]^{d,C} \subseteq D_\sigma$ and $[\![\gamma]\!]^{d,C} \subseteq D_{\langle\sigma,\tau\rangle}$, then $[\![\alpha]\!]^{d,C} = \{c(b) : b \in [\![\beta]\!]^{d,C} \wedge c \in [\![\gamma]\!]^{d,C}\}$
(Morzycki (2011), based on Kratzer & Shimoyama (2002))

To illustrate, suppose a function A, type $\langle e, st\rangle$ and a set B, type $\langle e\rangle$, as represented with the sets of alternatives in (25) below.

(25) $$A = \left\{ \begin{array}{l} \lambda x\lambda w.f(x)(w), \\ \lambda x\lambda w.g(x)(w), \\ \lambda x\lambda w.h(x)(w) \end{array} \right\}$$
$$B = \{a, b, c\}$$

Since these are sets, $A(B)$ has to proceed via HFA and not FA. Each object in the set A is applied to each object in B, resulting in $C = A(B)$. This is illustrated in (26).

(26) $$C = A(B) = \left\{ \begin{array}{lll} [\lambda x\lambda w.f(x)(w)](a), & [\lambda x\lambda w.f(x)(w)](b), & [\lambda x\lambda w.f(x)(w)](c), \\ [\lambda x\lambda w.g(x)(w)](a), & [\lambda x\lambda w.g(x)(w)](b), & [\lambda x\lambda w.g(x)(w)](c), \\ [\lambda x\lambda w.h(x)(w)](a), & [\lambda x\lambda w.h(x)(w)](b), & [\lambda x\lambda w.h(x)(w)](c) \end{array} \right\}$$
$$= \left\{ \begin{array}{l} \lambda w.f(a)(w), \lambda w.f(b)(w), \lambda w.f(c)(w), \\ \lambda w.g(a)(w), \lambda w.g(b)(w), \lambda w.g(c)(w), \\ \lambda w.h(a)(w), \lambda w.h(b)(w), \lambda w.h(c)(w) \end{array} \right\}$$

[3] Morzycki's analysis allows for parts of the grammar to crucially *not* be Hamblinized, but the details are not important here.

C here is the result of the pointwise function application of the elements from set A to set B. This results in the alternatives from both A and B being represented in C.

To summarize, the interpretation function is parameterized to a degree of precision d. This controls the size of a pragmatic halo of alternatives, generated with the $\approx$ relation. The halo is a set of alternatives that resemble the core predicate to some degree. Being sets rather than functions, there also needs to be a new notion of function application: Hamblin Function Application. I assume a framework such as this in my analysis of *sorta*.

4. Analysis

4.1. Approximation with sorta

As discussed in section 2.1, *sorta* has the flavor of an approximator. Intuitively, the meaning of *sorta* represents some form of approximation. This is evident when we consider the paraphrases available for sentences that use *sorta*, as in (27) and (28). These paraphrases have in common the fact that they note some "closeness" to the predicate being modified.

(27) He sorta swam over to the boat.
"He did something like swimming."

(28) The soccer player sorta kicked the ball.
"The player did something close to but not quite kicking the ball."

How can we profitably think about closeness? The suggestion I make here is to think about *sorta* as allowing a speaker to get out of saying something that would be false. Consider the kinds of situations *sorta* is used in. Generally, when a speaker makes a *sorta V* assertion, it is in a situation where the verb cannot be used very felicitously. *Sorta* allows a speaker to expand the meaning of the verb to encompass situations that it otherwise could not describe. For concreteness, take example (1) from earlier in the paper, repeated as (29) below. The speaker here is using *sorta kicked the ground* to describe how she hurt her foot, but presumably *sorta kicked the ground* was used because *kicked the ground* did not accurately describe the situation. *Sorta* here is being used to expand the meaning of *kick the ground* in order to accurately describe the situation.

(29) [I was] running on concrete and accidentally sorta kicked the ground.

The effect of *sorta* here should be compared with the behavior slack regulators such as *exactly*. Lasersohn's insight was that speakers allow each other an amount of pragmatic slack—false utterances can be pragmatically licit (and not uncooperative) if speakers allow each other to be imprecise. *Exactly* removes pragmatic slack, forcing speakers to be more precise. In other words, *exactly* allows for fewer expressions to count as "good enough" in a context. This is the opposite behavior of *sorta*. As Lasersohn notes, hedges like *sorta* expand the halo so that something in the halo is true.

In the framework assumed, halos are regulated through the imprecision parameter d on $[\![.]\!]$. Therefore, *sorta* needs to affect the degree of precision by existentially quantifying over a degree and setting that degree as the value for the imprecision parameter on the object that *sorta*

combines with (denoted as α in (30) below). By using the PREC typeshift, *sorta* can have access to the imprecision parameter, as PREC turns non-gradable predicates into gradable predicates by binding the imprecision parameter with a lambda.

Naturally, there need to be some constraints on the degree *sorta* introduces; *Con* is a placeholder describing these constraints. As the expression *sorta* combines with a set of alternatives (the pragmatic halo), *sorta* will need to pick something from this set. This is accomplished in the second conjunct in (30) below, with existential quantification over functions $f \in [\![\text{PREC } \alpha]\!]^{d'}(d)$ (i.e., functions in the halo of α after the imprecision parameter has been set). I assume that VPs denote properties of individuals, and so the function picked from α's alternatives will need to be applied to an individual. These elements are put together in the denotation in (30).

(30) (Tentative)
$[\![\textit{sorta } \alpha]\!]^{d'} = \lambda x \exists d \left[\textit{Con}(d) \wedge \exists f \in \left[[\![\text{PREC } \alpha]\!]^{d'}(d)\right] [f(x)]\right]$
where $[\![\alpha]\!]$ is type $\langle d, et\rangle$

Syntactically, I assume *sorta* is a VP adjunct, as in (31). In the case of this example, α would be the VP PREC *swim*. I ignore the category of *sorta* and labeled it XP.

(31)

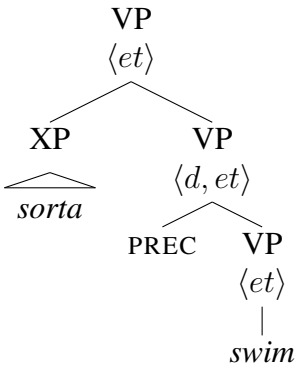

Superficially, $[\![\textit{sorta}]\!]$ behaves a function of type $\langle\langle d, et\rangle, et\rangle$, as it combines with a gradable predicate and an individual. Strictly speaking, however, *sorta* does not have this type — the gradable predicate is not an argument of *sorta*, but is introduced syncategormatically. Writing *sorta* this way reflects how Morzycki (2011) and Kratzer & Shimoyama (2002) develop their semantics, but it also is a notational convenience. See Rawlins (2008) for thoughts on how to redevelop aspects of Kratzer & Shimoyama (2002) using functions rather than sets.

The degree introduced by *sorta* is used as the degree of precision for the expression *sorta* modifies. We need to return to the constraints on this degree, *Con*. The degree of precision controls the size of the pragmatic halo associated with some expression. As *sorta* is analyzed as expanding a pragmatic halo, the question here is how much the halo should be expanded. The halo should be expanded only a little amount; too much halo expansion, and an expression modified by *sorta* could come to mean anything. That *sorta* has an approximative meaning — *sorta kick* intuitively has a meaning close to *kick* — tells us that *sorta* expands the halo to only some small degree. To capture this, I introduce an operator $\lessdot$ that is true just in case one degree is less than but close to the value of a second degree. This is defined in (32).

(32) $\forall d \forall d'$, $d \lessdot_C d'$ iff $d < d'$ and the value of d is close to d' as determined by the context C.

I assume that the new degree of precision is close to the standard degree of precision for the context — the degree of precision that the interlocutors have (usually implicitly) agreed to use when determining whether an utterance is true or false. This is comparable to the notion of a standard in the semantics of gradable adjectives. For instance, relative adjectives such as *tall* are often analyzed as being associated with standards that determine whether the adjective holds of some entity (Kennedy 2007; Kennedy & McNally 2005; Bierwisch 1989; Cresswell 1976). The degree of precision used in a discourse can shift, as slack regulators show, and standards associated with relative adjectives can also shift (the standard for whether a person is tall is not the same standard for whether a building is tall, for instance). The standards associated with verbs also change; what counts as a kick in one context might not necessarily count as a kick in a different context (an infant kicking a ball compared to a professional soccer player kicking a ball, perhaps).

Following Kennedy (2007), I assume a function **standard** that maps a gradable predicate (type $\langle d, et\rangle$) to the degree necessary for that predicate to hold true in the context. As the analysis pursued here crucially relies on a standard degree of precision, I assume that **standard** can return the standard degree of precision for a verbal predicate in addition to its typical duty with adjectives. This requires no change in the types that **standard** is defined over, since $[\![\text{PREC } \alpha]\!]$ (where α is a verb phrase) will be type $\langle d, et\rangle$.

The content of *Con*, then, is to compare the new degree of precision with the standard degree of precision using $\lessdot$. The denotation for *sorta* is updated in (33) to reflect this. (I have suppressed the context argument on $\lessdot$ for presentational clarity.)

(33) (Final)
$$[\![\textit{sorta } \alpha]\!]^{d',C} = \lambda x \exists d \left[d \lessdot \textbf{standard}([\![\alpha]\!]) \wedge \exists f \in \left[[\![\alpha]\!]^{d'}(d) \right] [f(x)] \right]$$

To demonstrate, *sorta swim* would be translated as (34) below. *sorta* has combined with the VP PREC *swim*.

(34)
$$[\![\textit{sorta } \text{PREC } \textit{swim}]\!]^{d',C} = \lambda x \exists d \left[\begin{array}{l} d \lessdot \textbf{standard}([\![\text{PREC } \textit{swim}]\!]^{d',C}) \wedge \\ \exists f \in [[\![\text{PREC } \textit{swim}]\!]^{d',C}(d)][f(x)] \end{array} \right]$$

The imprecision parameter on $[\![\text{PREC } \textit{swim}]\!]^{d}$ is set to d. (35) is equivalent to (34), but PREC *swim* is rewritten using the $\approx$ notation.

(35)
$$[\![\textit{sorta } \text{PREC } \textit{swim}]\!]^{d',C} = \lambda x \exists d \left[\begin{array}{l} d \lessdot \textbf{standard}([\![\text{PREC } \textit{swim}]\!]^{d',C}) \wedge \\ \exists f \in \{f_{\langle e,t\rangle} : f \approx_{d,C} \textbf{swim}\}[f(x)] \end{array} \right]$$

Although the alternatives for any particular expression are context-dependent, for concreteness (34) might look as in (36), where the alternatives $\{\textbf{swim}, \textbf{float}, \textbf{wade}, \ldots\}$ are represented.[4]

[4] I simplify the alternatives here, but it should be assumed that, e.g., $\textbf{swim} = \lambda x.\textbf{swim}_w(x)$, with an open world variable.

(36) $[\![\textit{sorta}\ \text{PREC}\ \textit{swim}]\!]^{d',C}$

$$= \lambda x \exists d \left[\begin{array}{c} d \lessdot \textbf{standard}([\![\text{PREC}\ \textit{swim}]\!]^{d',C}) \wedge \\ \exists f \in \left\{ \begin{array}{c} \textbf{swim}, \\ \textbf{float}, \\ \textbf{wade}, \\ \ldots \end{array} \right\} [f(x)] \end{array} \right]$$

Worth reflecting on at this point is *sorta*'s similarity to another well-known morpheme: POS. It has been argued that in English and other languages, a phonologically null morpheme POS takes a gradable adjective as its argument (Cresswell 1976; von Stechow 1984; Bierwisch 1989; Kennedy 1999, 2007). The role of POS is to saturate the degree argument of the adjective and assert that the degree to which the entity holds the property denoted by the adjective meets or exceeds a contextually supplied standard.[5] POS can be stated as in (37) below. The similarity to note here is that both *sorta* and POS involve comparing a degree to a standard. This makes *sorta* a cousin to POS; instead of asserting that the standard is met, however, it asserts closeness to the standard.

(37) $[\![\text{POS}]\!] = \lambda G \lambda x \exists d\, [d \geq \textbf{standard}(G) \wedge G(d)(x)]$

In summary, *sorta* expands the pragmatic halo of some linguistic expression and picks a function from the expanded halo. The halo is expanded by using a degree of precision less than but near the standard degree of precision. This was accomplished through a new operator, $\lessdot$, and by generalizing the **standard** function to be able to pick out standard degrees of precision.

4.2. Hedging objects

In the previous section, I developed an analysis of how the verb may be hedged. The analysis depends on linguistic expressions having sets of alternatives available, alternatives that model Lasersohn's pragmatic halos. For instance, for a verb such as *swim*, each alternative is a function that resembles the core meaning of *swim*, the function **swim**, to some degree. The entire set of alternatives is a set of resembling alternatives that is ordered by their degree of resemblance to some function. What *sorta* does in this case is to lower the degree needed to be part of the set of resembling alternatives, by manipulating a degree of precision on the interpretation function.

But, as described earlier, *sorta* can also hedge the direct objects of some predicates. The question is how to get this kind of behavior with *sorta*, how *sorta* can hedge a direct object even when it does not merge with the direct object. The answer, I suggest, comes from the architecture of the Hamblin semantics assumed here.

In this system, denotations are represented as sets of alternatives that grow or shrink depending on the degree of precision. In the previous section, verbs (and verb phrases) were represented in this fashion, but we should expect that nouns (and noun phrases) are represented in this way as well, and this is precisely how Morzycki (2011) handles metalinguistic comparatives with

[5]This is a description of POS on the assumption that adjectives denote relations between degrees and individuals. See Kennedy (1999) for a different analysis of POS where adjectives denote measure functions.

nominals rather than adjectives. The denotation for *house*, for instance, would be represented as in (38), a set of functions f such that each resembles **house** to degree d.

(38) $[\![\textit{house}]\!]^d = \{f : f \approx_{d,C} \textbf{house}\}$

If nominals also have sets of resemblance alternatives, the problem of how *sorta* can hedge a direct object becomes the problem of how *sorta* has access to these alternatives. Put simply, the alternatives for the direct object have to "project" up to the VP level in order to be visible to *sorta*. The mechanism to do this is already available using an alternative semantics.

Since denotations are sets rather than functions, we required a new way of combining a predicate with its argument. This was Hamblin Function Application (HFA). The intuition formalized as HFA is to apply each function from one set of alternatives pointwise to its arguments in another set of alternatives. This creates a new set of alternatives with the alternatives of both the predicate and its argument. For concreteness, $[\![\textit{build a house}]\!]^d$ would be represented as in (39), with the alternatives from $[\![\textit{build}]\!]^d$ applying pointwise to the alternatives for $[\![\textit{house}]\!]^d$.[6]

(39) $[\![\textit{build a house}]\!]^d = \{b(h) : h \in [\![\textit{house}]\!]^d \wedge b \in [\![\textit{build}]\!]^d\}$

Since $[\![\textit{build a house}]\!]^d$ will have the alternatives of both $[\![\textit{build}]\!]^d$ and $[\![\textit{house}]\!]^d$, this solves the issue of how *sorta* has access to the alternatives of *house*. Quite simply, the alternatives from *house* will continue to project upwards to the VP level. Through this compositional process, HFA, the alternatives at a lower node in the syntax can project to higher nodes in the syntax. Abstracting away from the precise translation of *build* and *house*, the alternatives for *build a house* might project as in (40).

(40)

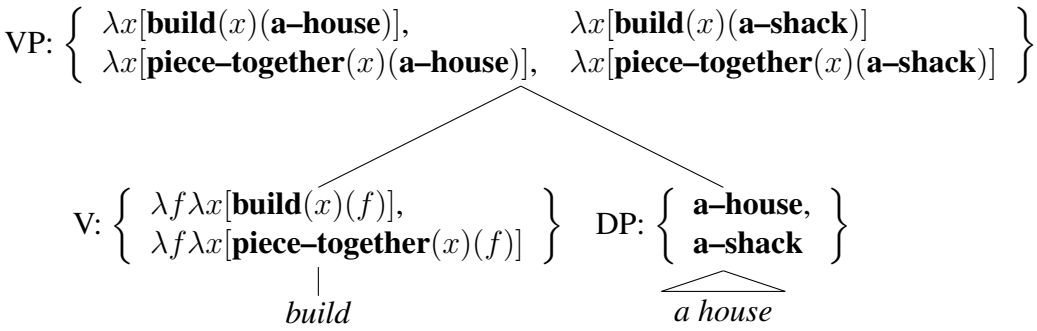

To summarize, *sorta* can modify the direct objects of verbs even when it has not merged with the direct object due to the mechanics of a Hamblin semantics. The reason for this comes from the behavior of Hamblin Function Application. HFA applies predicates from one set pointwise to arguments in a second set, creating a third set. This set contains all the alternatives from the first and the second set; in essence, HFA allows the alternatives from the direct object to percolate upward throughout the course of the derivation. Worth reflecting on here is that this behavior comes for free, since HFA is independently necessary in this framework. All things being equal, if alternatives are grammatically represented and certain expressions are sensitive to alternatives, we should expect cases of apparent non-local relationships between some expressions and alternative sensitive elements. Finding that *sorta* exhibits this behavior (albeit in limited ways) is less surprising in light of the alternative semantics I have adopted.

[6] I assume that the singular indefinite article *a* has no semantic contribution here, so that $[\![\textit{house}]\!]^d = [\![\textit{a house}]\!]^d$.

4.3. Summary and lingering issues

The previous two sections attempted a first pass at an analysis of the facts presented in section 2.1. This analysis relied on adopting the alternative semantics for imprecision proposed by Morzycki (2011). The behavior of *sorta* was to widen the pragmatic halo of an expression (modeled as a set of alternatives) and choose some function from this halo. This function resembled the predicate modified by *sorta* to some degree determined by a degree of precision parameter. By lowering the degree of precision, *sorta* could widen the halo. The fact that *sorta* could hedge the direct object of a verb and not simply the verb was a consequence of adopting an alternative semantics; Hamblin Function Application, necessary to make this sort of system compositional, allows for alternatives to project upwards throughout the course of a derivation. With the alternatives of both the verb and its object projecting up to the VP level, *sorta* was in a position to hedge both. But, there remain some issues with the analysis so far.

The first issue is accounting for the projection behavior of different sorts of noun phrases (indefinite versus definite). As repeated in the contrast in (41) and (42), *sorta* can hedge a singular indefinite direct object but not a definite direct object. Obviously, there must be some reason for this restriction.

(41) The carpenter sorta built a barn.
   a. The carpenter did something that was like building (e.g., putting together a prefabricated structure).
   b. The carpenter built something like a barn (e.g., a shed).

(42) The amateur carpenter sorta built the house.
   a. The amateur carpenter did something that was akin to building that resulted in the house (e.g., he had help from more skilled carpenters).
   b. *The amateur carpenter built something that was like the house (a shack, a hovel, . . .).

The second issue is related to the content of the verb itself. While only the alternatives for singular indefinites can project up to VP, most verbs seem to block the projection of the object's alternatives. However, some verbs, particularly verbs of creation and some intensional verbs, are holes with respect to the projection of alternatives. The account so far predicts that the alternatives of any indefinite singular direct object should project, which simply is not the case. In the next section, I develop explanations for these facts about the projection of resemblance alternatives.

5. Constraints on sorta

5.1. Determiner effects

As demonstrated earlier, the definite determiner blocks alternatives from projecting. The questions to pursue here are how alternatives are blocked, and why it should be the case that they are blocked.

The simplest way to block resemblance alternatives from projecting is by transforming the set of alternatives into a singleton set. HFA would proceed as before, but having only a single item in the set, that alternative would be the only alternative to project. There are two obvious possibilities for creating a singleton. One option is to simply specify maximum precision on the nominal, as in (43). This relies on $\approx$ being equivalent to $=$ when the degree parameter on $\approx$ is set to the maximal degree.

(43) $[\![\textit{house}]\!]^{d=1} = \{f : f \approx_d \textbf{house}\} = \{\textbf{house}\}$

This option loses its appeal, however, when we consider the fact that in typical speech we allow some slack in word choice, even if it is not signaled with *sorta*. Enforcing maximum precision would make the incorrect prediction that definites are always interpreted maximally precise. I suggest another option: the use of a choice function, mapping the set of alternatives to a single alternative. I make this part of the meaning of the definite determiner, as in (44).

(44) $[\![\textit{the } \alpha]\!]^d = \{\iota x.\textbf{choice}([\![\alpha]\!])(x)\}$

(45) $[\![\textit{the house}]\!]^d = \{\iota x.\textbf{choice}([\![\textit{house}]\!])(x)\}$

Speculating on why the definite might behave like this, one possible explanation is related to the fact that the definite presupposes the existence of individuals that satisfy the predicate. But, in the model, these individuals will already be true of the nominal predicate. Picking any alternatives besides the alternative the individual is true of would be false. This forces only a single alternative to project.

5.2. Verb class

In section 4.2, I provided an analysis of *sorta* to account for not only how hedging can occur, but also why the object of some verbs can be hedged. The answer, I suggest, relies on *sorta* behaving like a Lasersohnian slack regulator, widening the halo around a verb in order to include in the denotation of the verb things that might not otherwise "count" as part of the denotation. This in turn occurs by lowering the degree of precision required for interpretation, by setting the imprecision parameter on the interpretation function lower than the contextually supplied standard. Doing this increases the amount of imprecision alternatives available. Objects can be hedged in this system because the alternatives of the object can project to the VP level by combining with the alternatives of the verb pointwise.

This account still severely overgenerates on the readings possible. Namely, the account so far predicts that all objects should be able to be hedged. This in fact is not the case; hedging is severely constrained. Only some verbs allow for their objects to be hedged, and among those objects that can be hedged, it is only indefinite noun phrases and not definite noun phrases. The goal here is to provide an account of this, and constrain the system to allow only the attested readings.

To pursue an explanation here, I want to start by asking the question of what makes the verbs that allow for hedging of their object special. Creation verbs and intensional transitive verbs such as *look for* allowed for hedging. Are these special in any way? The literature on verbs has suggested that they are in fact special with respect to their direct object position.

The relevant feature here is whether this is an opaque argument position or not. It is well-known that intensional verbs such as *seek* and *look for* have a referentially opaque direct object position (Van Geenhoven & McNally 2005; Zimmermann 1993; Montague 1974; Quine 1964; Moltmann 1997). There is no entailment that an entity instantiating the property denoted by the indefinite exists.

(46) The man was looking for a horse.
   a. Transparent reading: There is a horse than the man was looking for.
   b. Opaque reading: The man is looking for a horse and it may or may not exist.

Creation verbs also exhibit this failure of existential quantification, as noted by von Stechow (2001). Holding the reference time constant, the argument in (47) does not hold. In comparison, a non-intensional, non-creation verb such as *push* allows this argument to go through (48). The reasoning for this plain; creation verbs only entail the existence of the created object at the end of the event.

(47) John drew a circle.
DOES NOT ENTAIL: There was a circle that John drew.

(48) John pushed a cart.
ENTAILS: There was a cart that John pushed.

Zimmermann (1993) and Van Geenhoven & McNally (2005) argue that intensional transitive verbs are special because of their argument structure. Intensional transitive verbs involve an attitude towards a property — they have a property-type argument — while non-intensional transitives have more mundane individual-type arguments. De Swart (2001) also argues that these verbs are special; intensional verbs allow for weak readings of indefinite noun phrases because these noun phrases have well-formed property-type denotations. These arguments are built on the referentiality of the noun phrase in object position; noticing the similarities between intensional transitive verbs and creation verbs with respect to their object position and existential exportation, we might extend this analysis to creation verbs and suggest that they also take property-type objects. What this amounts to, in the lexical semantics for these verbs, is local existential quantification over entities instantiating the property.

The data from *sorta* suggests that this is on the right track. Looking at intensional verbs, the natural reading for a noun phrase hedged by *sorta* as in (49) is one where existential exportation does not hold.

(49) He was sorta looking for a horse.
*'There is something like a horse that he is sorta looking for.'

I take the constraining factor here to be one of argument types. Verbs which allow for property-type arguments (type $\langle et, et\rangle$) allow for the alternatives of their direct object to project. Verbs which only allow for individual-type arguments ($\langle e, et\rangle$) do not allow the alternatives to project. This is a nod to claims that certain types of verbs are special with respect to the type of the arguments they combine with. Although I will not pursue a full formal analysis here, the hypothesis is that verbs that accept property-type complements allow for a sort of "escape hatch" through which the resemblance alternatives can project.

6. Discussion

6.1. Where resemblance alternatives come from

A relevant question to ask is where the alternatives for any particular expression come from. Formally, sets of alternatives were built using the $\approx$ relation, which was true just in case two functions resembled each other to some degree. But, this pushes back the explanation on where alternatives come from to the mechanics of $\approx$. A real explanation for why we see the alternatives we do would be ideal. Put more concretely, why should the alternatives to *swim* include *wade*, *float*, and so on?

A point to note here is that the alternatives available do not have to be represented by lexical items themselves. For instance, *sorta kick the ground* expresses something like a kicking action, but the purpose of the speaker using *sorta* here seems to be to show that the conceptual content of *kick* does not precisely match what happened in the event. This contrasts with a focus-sensitive adverb such as *only*; the alternatives *only* invokes, although they depend on the meaning of the focused element, are most naturally associated with other lexical items.

(50) a. Suzy sorta jogged.
b. Suzy only jogged$_F$. (not ran or sprinted)

(51) a. John sorta swam to the boat.
b. John only swam$_F$ to the boat. (but didn't climb in to the boat)

The alternatives *sorta* gets access to seem to be related in some way to the conceptual content of the lexical item that is being hedged. The alternatives invoked with *sorta swim*, for instance, although they might not be called swimming, intuitively are associated with swimming in some way: moving the arms and legs, floating in water, and so on.

Two points should be noted, then. First, $\approx$ has access to the conceptual content of the functions that it compares, in order to determine whether two functions resemble each other to the required degree. More work might be done on fleshing out $\approx$, perhaps in terms of prototypes (Kamp & Partee 1995), and that $\approx$ (in a pretheoretic sense) is a window into the interface between formal semantics and the conceptual system. Second, the alternatives that we see do not have to named by lexical items. This seems to requires that the functional domains in the semantic model have an infinite (and perhaps dense) space of functions, as the building of the set of alternatives requires comparison to any number of functions. Lexical items carve up this space in vague, context-dependent ways, depending on the precision required, but there are gaps not covered by particular lexical items in normal circumstances that *sorta* gives the speaker access to.

6.2. Restriction and the projection of alternatives

I provide an analysis of why objects can be hedged in section 4.2 and attempt to explain the restrictions on this in section 5. However, I suggest here a second analytical option for why objects can be hedged and those constraints. This analysis is based on the work of Chung & Ladusaw (2004). Here, what I suggest is that intensional transitive verbs and creation verbs

combine with their objects differently than other verbs. Namely, these verbs will combine via Chung and Ladusaw's Restrict mode of composition, while other verbs combine via Specify. A difference in the mode of composition correlates with whether the alternatives for the direct object are visible to *sorta*.

Chung & Ladusaw (2004), based on data from Chamorro and Maori, argue for two new modes of semantic composition, what they call Restrict and Specify. Part of the problem they are trying to solve in introducing new modes of composition is the behavior of indefinites. Indefinites have a range of behaviors associated with them. Some authors have argued that this calls for a flexible type system where indefinites can be translated between quantificational, property, and individual types (Partee 1987). The approach Chung and Ladusaw argue for is that indefinites have uniformly property-type denotations, but that there exist different modes of semantic composition with different semantic effects.

They introduce modes of composition they call Restrict and Specify. Restrict contrasts with the familiar Function Application by being a non-saturating mode of composition, leaving a lambda untouched in the derivation; Function Application is a saturating mode of composition. The effect of this is illustrated in the hypothetical example in (52) (the derivation is impossible in English). Here, $[\![cat]\!]$ Restricts λy in the denotation of $[\![bit]\!]$. Conceptually, this is a form of intersection, intersecting cats with things that were bit.

(52) The dog bit cat.

a. **Restrict**$(\lambda x[\textbf{cat}(x)], \lambda y \lambda z[\textbf{bit}(z)(y)])$

b. $\lambda y \lambda z[\textbf{bit}(z)(y) \wedge \textbf{cat}(y)]$ (via Restrict)

c. $\lambda z \exists y[\textbf{bit}(z)(y) \wedge \textbf{cat}(y)]$ (via Existential Closure)

Because Restrict is non-saturating, predicates still require some way of being saturated. Function Application with the open argument position is one way. Existential Closure is a second way (illustrated in (52c). They assume that all unsaturated predicates undergo existential closure of their open arguments at what they call the event level (roughly corresponding syntactically to *v*P). This has the effect of making it so that arguments composed via Restrict take obligatory narrow-scope with respect to negation and other operators.

The second mode of composition they introduce is Specify. Specify involves a local type-shift — a choice function. Choice functions map properties to entities, so the choice function can be an argument to a predicate that is looking for an entity-type argument. Choice functions introduce a semantic unfulfilledness — the choice function needs to be bound by an existential somewhere in the derivation — but the predicate can be saturated with a choice function. Existential closure over the choice function can happen at the event level or the clausal level (or both), depending on the parameters of the language. Specify therefore allows for both wide-scope and narrow-scope indefinites. I illustrate Specify in (53).

(53) The dog bit a cat.

a. **Specify**$(f(\textbf{cat}), \lambda x \lambda y[\textbf{bit}(y)(x)]$

b. $\lambda y[\textbf{bit}(y)(f(\textbf{cat}))]$ (via CF applied to **cat**, FA)

c. $\lambda y \exists f[\textbf{bit}(y)(f(\textbf{cat}))]$ (via EC)

Indefinites in intensional transitive verbs and creation verbs take low-scope with respect to negation. For instance, *not build a house* does not have the inference that there exists a house that was not built. Rather, the inference that goes through is that no house was built. Under the system described, this could be interpreted as the indefinite combining with the verb via Restrict. The examples in (49) also suggest that this is the case. Taking a view of referential opacity as narrow-scope with respect to an intensional operator, these also suggest that the indefinite is taking narrow-scope. Again, this can be represented as the indefinite composing with the predicate via Restrict.

Indefinites in object position for other types of verbs, outside of the creation and intensional transitive verb classes discussed, do allow for both wide and narrow-scope readings. My interpretation of this fact is that indefinite objects for these verbs compose instead via Specify, with the existential closure over the choice function happening at either the event or the clausal level.

The picture that emerges here is that Restrict allows for the alternatives of indefinites to project, while Specify does not. The cases where I argued that Restrict was active are those cases where the alternatives of an indefinite object are accessible to *sorta*. I will suggest that the relevant difference between Specify and Restrict here is whether they are saturating modes of composition. Saturation closes off imprecision alternatives, while non-saturating modes of composition do not.

This conclusion is supported by the behavior of definites. By their nature, definites have individual-type denotations, and therefore do not need a special mode of composition; they can compose with predicates by Function Application. Function Application is a saturating mode of composition, which would predict that definites can never be hedged. This is in fact the case; the alternatives for definite noun phrases are never accessible to *sorta*.

To summarize the idea here, saturating modes of predication close off sets of imprecision alternatives, while non-saturating modes of predication do not. Creation verbs and intensional verbs combine with indefinite objects via Restrict, a non-saturating mode of composition, and so allow for their objects to be hedged by *sorta*. I will not pursue an explanation here for why saturation versus non-saturation matters for the projection of alternatives, but framing the problem in this way provides another point of entry into explaining the projection behavior of resemblance alternatives.

7. Conclusion

In this paper I have presented an account of *sorta*. My analysis depended on building up sets of resemblance alternatives, with *sorta* picking a single alternative from a set. This set of alternatives was meant to model pragmatic halos (Lasersohn 1999; Morzycki 2011). Much work remains to be done in expanding the range of the analysis, however. This was admitted as such in section 5, where I attempted to show how to constrain the readings available to *sorta* with certain verbs and indefinite objects, linking the projection of alternatives up to whether a verb accepts a property-type complement and whether the determiner on the DP direct object presupposes its domain. I discuss these facts, and provide the outline of an alternate solution using Chung & Ladusaw (2004)'s Restrict and Specify modes of composition. Finally, I discussed some consequences this work on *sorta* has for the semantics and conceptual system interface,

suggesting that the ≈ relation, which generates resemblance alternatives, has access to the conceptual content of a predicate, and that the alternatives generated do not need to be represented by any particular lexical item.

Acknowledgements

I thank Marcin Morzycki, Cristina Schmitt, Alan Munn, Alan Beretta, Karthik Durvasula, the Michigan State University semanticists, Greg Johnson, Jessica Gamache, Hannah Forsythe, and an anonymous ConSOLE proceedings reviewer for their helpful comments and suggestions, as well as audiences at ConSOLE XXI, Penn Linguistics Colloquium 37, and SALT 23. All errors are my own.

Curt Anderson
Michigan State University
ande1472@msu.edu
http://www.msu.edu/~ande1472

References

Abney, S. (1987). *The English noun phrase in its sentential aspect*. PhD thesis, MIT.
Bierwisch, M. (1989). The semantics of gradation. Bierwisch, M. & E. Lang (eds.), *Dimensional Adjectives*, Springer, Berlin, pp. 71–261.
Bolinger, D. (1972). *Degree Words*. Mouton, The Hague.
Chung, S. & W. Ladusaw (2004). *Restriction and saturation*. MIT Press.
Cresswell, M. J. (1976). The semantics of degree. Partee, B. (ed.), *Montague Grammar*, Academic Press, New York, pp. 261–292.
Giannakidou, A. & M. Stavrou (2009). Metalinguistic comparatives and negation in Greek. Halpert, C., J. Hartman & D. Hill (eds.), *Proceedings of the 2007 Workshop on Greek Syntax and Semantics*.
Giannakidou, A. & S. Yoon (2011). The subjective mode of comparison: Metalinguistic comparatives in Greek and Korean. *Natural Language & Linguistic Theory* 29:3, pp. 621–655.
Grice, H. (1957). Meaning. *The philosophical review* pp. 377–388.
Hamblin, C. (1973). Questions in Montague English. *Foundations of Language* 10:1, pp. 41–53.
Heim, I. & A. Kratzer (1998). *Semantics in generative grammar*. Wiley-Blackwell.
Kamp, H. & B. Partee (1995). Prototype theory and compositionality. *Cognition* 57:2, pp. 129–191.
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, pp. 1–45.
Kennedy, C. & L. McNally (2005). Scale structure, degree modification, and the semantics of gradable predicates. *Language* 81:2, pp. 345–381.
Kratzer, A. & J. Shimoyama (2002). Indeterminate pronouns: The view from Japanese. Otsu, Y. (ed.), *Third Tokyo Conference on Psycholinguistics*, Hituzi Syobo, Tokyo, pp. 1–25.
Lasersohn, P. (1999). Pragmatic halos. *Language* 75:3, pp. 522–551.
McCawley, J. D. (1998). *The Syntactic Phenomena of English*. University of Chicago Press.
Moltmann, F. (1997). Intensional verbs and quantifiers. *Natural Language Semantics* 5:1, pp. 1–52.
Montague, R. (1974). *Formal philosophy*. Yale University Press.

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

Partee, B. (1987). Noun phrase interpretation and type-shifting principles. *Studies in discourse representation theory and the theory of generalized quantifiers* 8, pp. 115–143.

Quine, W. (1964). *Word and object*, vol. 4. MIT press.

Rawlins, K. (2008). *(Un)conditionals: an investigation in the syntax and semantics of conditional structures*. PhD Thesis, UCSC.

Rooth, M. (1985). *Association with focus*. PhD Thesis, University of Massachusetts.

Rooth, M. (1992). A theory of focus interpretation. *Natural language semantics* 1:1, pp. 75–116.

Siegel, M. (2002). *Like*: The discourse particle and semantics. *Journal of Semantics* 19:1.

de Swart, H. (2001). Weak readings of indefinites: type-shifting and closure. *The Linguistic Review* 18:1, pp. 69–96.

Van Geenhoven, V. & L. McNally (2005). On the property analysis of opaque complements. *Lingua* 115:6, pp. 885–914.

von Stechow, A. (1984). Comparing semantic theories of comparison. *Journal of Semantics* 3:1, pp. 1–77.

von Stechow, A. (2001). Temporally opaque arguments in verbs of creation. *Semantic Interfaces: Reference, Anaphora, and Aspect*, CSLI Publications, pp. 278–319.

Zimmermann, T. (1993). On the proper treatment of opacity in certain verbs. *Natural Language Semantics* 1:2, pp. 149–179.